

GL₂ ℝ ORBIT CLOSURES IN HYPERELLIPTIC COMPONENTS OF STRATA

PAUL APISA

1. INTRODUCTION

The dynamics of the GL₂ ℝ action on strata of abelian differentials strikingly parallels the dynamics of unipotently generated flows on lattice quotients of semisimple Lie groups. In the homogeneous setting Ratner's theorem implies that orbit closures are lattice quotients of Lie subgroups. In the Teichmüller dynamics setting, Eskin, Mirzakhani [6] and Eskin, Mirzakhani, and Mohammadi [7] imply that orbit closures are affine invariant submanifolds, i.e. submanifolds cut out by real linear homogeneous equations in period coordinates. A classification of these orbit closures remains elusive.

The first progress made on this question was by McMullen [15] who showed that all orbit closures in $\mathcal{H}(2)$ are either $\mathcal{H}(2)$ itself or closed orbits, all of which had been discovered by Calta [5] and McMullen [13] and whose components were classified by McMullen [13]. In [15], McMullen classified orbit closures in $\mathcal{H}(1, 1)$ showing that they were either closed orbits, of which there is just one (see McMullen [14]), three-dimensional loci of eigenforms of real multiplication, or $\mathcal{H}(1, 1)$. These results were achieved a decade prior to Eskin, Mirzakhani [6] and Eskin, Mirzakhani, and Mohammadi [7] and establish analogues of those theorems in the genus two setting.

Motivating Question 1. *Classify orbit closures in every component of every stratum of abelian differentials.*

Apart from examples in genus two very few geometrically primitive orbit closures are known. An orbit closure \mathcal{M} fails to be geometrically primitive if for every abelian differential $(X, \omega) \in \mathcal{M}$ - where X is a Riemann surface and ω a holomorphic one-form - there is a quadratic differential (Y, q) on a Riemann surface Y and a holomorphic map $f : X \rightarrow Y$ such that $\omega^2 = f^*q$. An orbit closure that is not geometrically primitive is called a branched covering construction. To explain the paucity of known orbit closures Mirzakhani articulated two conjectures, which first appeared in [21].

Conjecture (Mirzakhani Arithmeticity Conjecture). *Higher rank affine invariant submanifolds are arithmetic.*

Conjecture (Mirzakhani Covering Conjecture). *Higher rank arithmetic affine invariant submanifolds are either connected components of strata or branched covering constructions.*

While strong evidence suggests that the Mirzakhani conjectures are both false in general, we will see that the Mirzakhani conjectures remain true in at least two components of strata of abelian differentials in each genus $g \geq 2$.

Motivating Question 2. *In which components of which strata do the Mirzakhani conjectures hold?*

Kontsevich and Zorich classified the connected components of strata of abelian differentials in [8].

Theorem (Kontsevich-Zorich). *All strata are connected except for the following:*

- $\mathcal{H}(2g-2)$, which has three connected components characterized by hyperellipticity and even/odd spin when $g > 3$.
- $\mathcal{H}(g-1, g-1)$, which has three connected components characterized hyperellipticity and even/odd spin when g is odd.
- $\mathcal{H}(g-1, g-1)$, which has two connected components characterized by hyperellipticity and nonhyperellipticity when g is even.
- $\mathcal{H}(2k_1, \dots, 2k_n)$, which has two connected components characterized by even/odd spin when $g > 3$ (excluding the case $\mathcal{H}(g-1, g-1)$ which has already been covered).
- $\mathcal{H}(4)$ and $\mathcal{H}(2, 2)$, which have two connected components - a hyperelliptic and an odd one.

The hyperelliptic components are simplest to understand. The underlying curves in $\mathcal{H}^{hyp}(2g-2)$ (resp. $\mathcal{H}^{hyp}(g-1, g-1)$) have the form $y^2 = f(x)$ where f is a simple polynomial of degree $2g+1$ (resp. $2g+2$) and the holomorphic one-forms have the form $\frac{dx}{y}$.

The geometric description of hyperelliptic components of strata is similarly straightforward. Let X be a hyperelliptic Riemann surface with hyperelliptic involution J and let ω be a holomorphic one-form on X . If ω has a single zero at p then (X, ω) is contained in $\mathcal{H}^{hyp}(2g-2)$ if and only if p is a Weierstrass point. If ω has two zeros at p and q of equal order then (X, ω) is contained in $\mathcal{H}^{hyp}(g-1, g-1)$ if and only if $Jp = q$.

Motivating Question 3. *Given a hyperelliptic curve $y^2 = f(x)$ where f is a simple polynomial, what is the $\mathrm{GL}_2 \mathbb{R}$ orbit closure of $\frac{dx}{y}$?*

In this paper we will establish:

Main Theorem 1. *If \mathcal{M} is an affine invariant submanifold in $\mathcal{H}^{hyp}(2g-2)$ or $\mathcal{H}^{hyp}(g-1, g-1)$ and $\dim_{\mathbb{C}} \mathcal{M} > 3$ then \mathcal{M} is a branched covering construction.*

In particular, the Mirzakhani conjectures are true for hyperelliptic components of strata. The main result is actually more specific, it states:

Main Theorem 2. *Under the same hypotheses as Theorem 1, if $\dim \mathcal{M} = 2r$ then \mathcal{M} is a branched covering construction over $\mathcal{H}^{hyp}(2r-2)$ and if $\dim \mathcal{M} = 2r+1$ then \mathcal{M} is a branched covering construction over $\mathcal{H}^{hyp}(r-1, r-1)$. The covers are branched over the zeros of the holomorphic one-forms and commute with the hyperelliptic involution.*

1.1. Finiteness of Teichmüller Curves.

As the following results show the Mirzakhani conjectures imply finiteness results for primitive Teichmüller curves.

Proposition. *If the Mirzakhani conjectures hold in a component \mathcal{H} of a stratum of abelian differentials then either \mathcal{H} contains finitely many geometrically primitive rank one orbit closures or any sequence of geometrically primitive Teichmüller curves equidistributes in \mathcal{H} .*

Proof. Suppose not to a contradiction. Let C_i be an infinite sequence of distinct geometrically primitive Teichmüller curves. By Eskin-Mirzakhani [6] this sequence equidistributes in a finite union of connected affine invariant submanifolds $\mathcal{M} = \bigcup_i \mathcal{M}_i$. Suppose that \mathcal{M} does not coincide with \mathcal{H} . The Mirzakhani conjectures imply that no \mathcal{M}_i is higher rank since this would imply that C_i is not geometrically primitive. Moreover, no \mathcal{M}_i is rank one since these orbit closures only contain finitely many nonarithmetic Teichmüller curves by Laneeau-Nguyen-Wright [9]. \square

Conjecture (Nonarithmetic Rank One Sparsity Conjecture). *Outside of genus two, a sequence of nonarithmetic rank one orbit closures cannot equidistribute in a component of a stratum.*

Conditional on the sparsity conjecture there are finitely many geometrically primitive Teichmüller curves in $\mathcal{H}^{hyp}(2g-2)$ and $\mathcal{H}^{hyp}(g-1, g-1)$ for $g > 2$. The work of Matheus-Wright [12] allows us to derive unconditional finiteness results for algebraically primitive Teichmüller curves.

Main Theorem 3. *In $\mathcal{H}^{hyp}(2g-2)$ and $\mathcal{H}^{hyp}(g-1, g-1)$ there are finitely many algebraically primitive Teichmüller curves for $g > 2$.*

Proof. By Matheus-Wright [12] algebraically primitive Teichmüller curves cannot equidistribute in the connected component of any stratum when $g > 2$. Since the Mirzakhani conjectures hold the previous argument implies that there are only finitely many algebraically primitive Teichmüller curves \square

In the case of $\mathcal{H}^{hyp}(g-1, g-1)$ this result was shown by Möller [18].

1.2. Classification of Orbit Closures.

We now offer a coarse classification of orbit closures in hyperelliptic components of strata. By coarse classification we mean a classification of orbit closures up to finitely many nonarithmetic closed orbits and up to classifying connected components of orbit closures.

Main Theorem 4 (Classification of Orbit Closures in $\mathcal{H}^{hyp}(2g-2)$). *The affine invariant submanifolds in $\mathcal{H}^{hyp}(2g-2)$ are:*

- (1) *Countably many Teichmüller curves which arise from torus covers branched over one point.*
- (2) *(If $g \equiv 2 \pmod{3}$) Countably many Teichmüller curves whose trace fields are degree two and which equidistribute in an affine invariant submanifold of $\mathcal{H}^{hyp}(2g-2)$ arising from a branched cover construction over $\mathcal{H}(2)$.*
- (3) *Conditional on the sparsity conjecture, finitely many Teichmüller curves beyond the previous two families.*

- (4) *Finitely many rank $r > 1$ affine invariant submanifolds for each $2r - 1 \mid 2g - 1$. These are branched covering constructions over $\mathcal{H}(2r - 2)$.*

Corollary (Characterization of Optimal Dynamics in $\mathcal{H}^{hyp}(2g-2)$). *Every orbit in $\mathcal{H}^{hyp}(2g-2)$ is either closed or equidistributed if and only if $2g - 1$ is prime.*

Proof. Higher rank proper orbit closures arise if and only if $2r - 1$ divides $2g - 1$ for some $r > 1$. \square

In the case of $g = 2$ this characterization of optimal dynamics follows from McMullen's classification of orbit closures in genus two. In the case of $g = 3$ this theorem was the main theorem of Nguyen-Wright [19].

Main Theorem 5 (Classification of Orbit Closures in $\mathcal{H}^{hyp}(g - 1, g - 1)$). *The affine invariant submanifolds in $\mathcal{H}^{hyp}(g - 1, g - 1)$ are:*

- (1) *Countably many Teichmüller curves which arise from torus covers branched over one point.*
- (2) *(If $g \equiv 0 \pmod{3}$) Countably many Teichmüller curves whose trace fields are degree two and which equidistribute in affine invariant submanifolds that are branched covering constructions over $\mathcal{H}(2)$.*
- (3) *Conditional on the sparsity conjecture, finitely many Teichmüller curves beyond the previous two families.*
- (4) *Countably many orbit closures that are branched covering constructions over $\mathcal{H}(0, 0)$.*
- (5) *(If g is even) Countably many orbit closures that cover genus two eigenform loci; these equidistribute in affine invariant submanifolds that are branched covering constructions over $\mathcal{H}(1, 1)$.*
- (6) *Conditional on the sparsity conjecture, finitely many three dimensional orbit closures beyond the previous two families.*
- (7) *Finitely many rank $r > 1$ affine invariant submanifolds for each $r \mid g$. These are branched covering constructions of $\mathcal{H}^{hyp}(r - 1, r - 1)$.*
- (8) *Finitely many rank $r > 1$ affine invariant submanifolds for each $2r - 1 \mid g$. These are branched covering constructions of $\mathcal{H}^{hyp}(2r - 2)$.*

In the case of $g = 2$ this theorem follows from McMullen's classification of orbit closures in genus two. In the case of $g = 3$ this theorem was one of the main theorems of Aulicino-Nguyen [1]

Corollary. *There are no higher rank proper even dimensional orbit closures in $\mathcal{H}^{hyp}(g - 1, g - 1)$ if and only if $g = 2^n$ for some n . There are no higher rank proper odd dimensional orbit closures in $\mathcal{H}^{hyp}(g - 1, g - 1)$ if and only if g is prime.*

The sparsity conjecture implies that in $\mathcal{H}^{hyp}(2g - 2)$ and $\mathcal{H}^{hyp}(g - 1, g - 1)$ we have classified every orbit closure up to finitely many geometrically primitive rank one orbit closures per genus. The results of this paper represent the first classification of orbit closures in any component of any stratum outside of genus two and three.

1.3. Background.

The cotangent bundle of the moduli space \mathcal{M}_g of smooth genus g curves is naturally identified with the bundle of quadratic differentials over \mathcal{M}_g . Each quadratic differential on a Riemann surface X associates a natural flat structure - i.e. a metric that is flat away from finitely many cone points - to X , see Zorich [24]. This flat structure is called a half-translation surface structure since it endows X with an atlas of charts to \mathbb{C} with transition functions given by $z \rightarrow \pm z + c$ for some $c \in \mathbb{C}$. The $\mathrm{GL}_2 \mathbb{R}$ action on \mathbb{C} induces a $\mathrm{GL}_2 \mathbb{R}$ action on half-translation surfaces and hence a $\mathrm{GL}_2 \mathbb{R}$ action on the cotangent bundle of \mathcal{M}_g .

Teichmüller's theorem states that if X and Y are distinct Riemann surfaces in \mathcal{M}_g that are distance d apart in the Teichmüller metric then the geodesic from X to Y is given by fixing a quadratic differential q on X , associating the natural flat structure to (X, q) , and then applying the matrix $\begin{pmatrix} e^d & 0 \\ 0 & e^{-d} \end{pmatrix}$. Therefore the $g_t := \begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}$ action on the bundle of quadratic differentials is called the Teichmüller geodesic flow. It follows that the $\mathrm{GL}_2 \mathbb{R}$ action on the bundle of quadratic differentials is the smallest group action generated by complex scalar multiplication and Teichmüller geodesic flow.

The cotangent bundle of moduli space contains a subbundle $\Omega\mathcal{M}_g$ of quadratic differentials that are squares of abelian differentials. $\Omega\mathcal{M}_g$ is stratified by specifying the number of zeros and their degree of vanishing on the underlying one-form. Let \mathcal{H} be such a stratum and let $(X, \omega) \in \mathcal{H}$. Let S be a symplectic basis of relative homology $H_1(X; Z(\omega))$ where $Z(\omega)$ is the zero set of ω . Local coordinates around (X, ω) are given by the map $\Phi(Y, \eta) = \left(\int_s \eta \right)_{s \in S}$. These coordinates are called period coordinates.

Recent work of Eskin, Mirzakhani, and Mohammadi implies that $\mathrm{GL}_2 \mathbb{R}$ orbit closures in strata of abelian differentials are particularly nice:

Theorem (Eskin-Mirzakhani [6]; Eskin-Mirzakhani-Mohammadi [7]). *$\mathrm{GL}_2 \mathbb{R}$ orbit closures in strata of abelian differentials are affine invariant submanifolds, i.e. $\mathrm{GL}_2 \mathbb{R}$ -invariant manifolds that are locally cut out by real homogeneous linear equations in period coordinates.*

For a survey of this theorem and its applications to the study and classification of affine invariant submanifolds see Wright [23]. Wright [23] also provides an overview of many of the techniques used in this paper.

The tangent bundle of an affine invariant submanifold \mathcal{M} at a point (X, ω) is naturally identified with a complex linear subspace of $H^1(X, Z(\omega); \mathbb{C})$, see [22] for details. Let $p : H^1(X, Z(\omega); \mathbb{C}) \rightarrow H^1(X, \mathbb{C})$ be the projection from relative to absolute cohomology.

Theorem (Avila-Eskin-Möller [3]). *If \mathcal{M} is an affine invariant submanifold, then $p(T_{(X, \omega)}\mathcal{M})$ is a complex symplectic vector space.*

Define the rank of \mathcal{M} to be $\mathrm{rk}(\mathcal{M}) := \frac{1}{2} \dim_{\mathbb{C}} p(T_{(X, \omega)}\mathcal{M})$ and define $\mathrm{rel}(\mathcal{M}) := \dim_{\mathbb{C}} \mathcal{M} - 2 \cdot \mathrm{rk}(\mathcal{M})$. An affine invariant submanifold is said to be higher rank if its rank is larger than 1.

1.4. Survey of Other Finiteness and Classification Results. Aside from the classification in genus two and the results of this paper the following summarizes current knowledge

of finiteness results for Teichmüller curves and classification results for higher rank orbit closures:

- (1) There are finitely many algebraically primitive Teichmüller curves in genus three (Bainbridge-Habegger-Möller [4]).
- (2) There are finitely many algebraically primitive Teichmüller curves in $\mathcal{H}(2g - 2)$ when g is an odd prime (Matheus-Wright [12]).
- (3) There are finitely many Teichmüller curves in any nonarithmetic rank one orbit closure (Lanneau-Nguyen-Wright [9]).
- (4) The Mirzakhani conjectures hold in all strata with two or fewer zeros in genus three. Conditional on the sparsity conjecture all such strata have finitely many geometrically primitive Teichmüller curves (Nguyen-Wright [19], Aulicino-Nguyen-Wright [2], Aulicino-Nguyen [1]).

1.5. Organization of the Paper and Remarks on the Proof. In Section 2 we will discuss a combinatorial model - the Lindsey half-tree - created by Kathryn Lindsey in [10] to study horizontally periodic translation surfaces in hyperelliptic components of strata. This model explains much of why orbit closures are particularly well-behaved in hyperelliptic components of strata. In Section 3 we will define branched covering constructions rigorously and devise a criterion for when an affine invariant submanifold is a branched covering construction. In Section 4 we will discuss Alex Wright's cylinder deformation theorem [22] and related constructions. In Section 6 we discuss Maryam Mirzakhani and Alex Wright's translation surface degeneration theorem [17]. In Section 7 we specialize these results to the setting of hyperelliptic components of strata. These sections establish the tools needed to run the basic mechanism of the proof: find a horizontally periodic translation surface in an affine invariant submanifold \mathcal{M} , use the cylinder deformation theorem to degenerate it to the boundary of \mathcal{M} , use the results of Section 2 to show that the boundary translation surface is a disjoint union of translation surfaces in hyperelliptic components of strata, and then use induction and the degeneration theorem to say something about the original translation surface.

In Section 5 we kick off the induction argument by establishing a host of nice properties satisfied by odd dimensional orbit closures in $\mathcal{H}^{hyp}(g - 1, g - 1)$. In Section 8 we establish criteria for affine invariant submanifold in hyperelliptic components of strata to be branched covering constructions. In this section we show that every rank $r > 1$ odd dimensional affine invariant submanifold in a hyperelliptic component of a stratum is a branched covering construction of $\mathcal{H}^{hyp}(r - 1, r - 1)$. In the final section, Section 9, we conclude by showing that every rank $r > 1$ even dimensional affine invariant submanifold in a hyperelliptic component of a stratum is a branched covering construction of $\mathcal{H}^{hyp}(2r - 2)$.

ACKNOWLEDGEMENT The author gratefully acknowledges research support by the National Science Foundation through the Graduate Research Fellowship Program. He thanks Alex Eskin, Alex Wright, and David Aulicino for their insightful and extensive comments. He thanks Kathryn Lindsey, Martin Möller, Anton Zorich, and Elise Goujard for helpful conversations.

2. THE LINDSEY HALF-TREE OF HORIZONTALLY PERIODIC TRANSLATION SURFACES IN HYPERELLIPTIC COMPONENTS OF STRATA

A half-graph Γ consists of a set of vertices, a set of edges (each of which connects two vertices), and a set of half-edges that begin at a vertex but do not end at a vertex. A half-tree is a half-graph whose vertex and edge set form a tree.

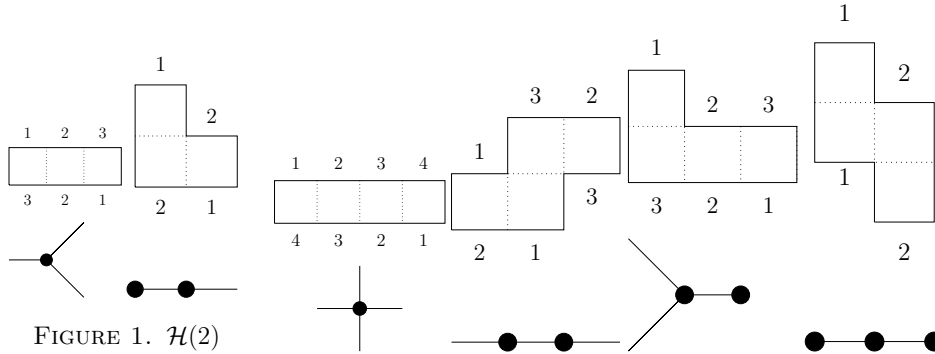
Theorem (Lindsey [10]). *A horizontally periodic translation surface in a hyperelliptic connected component of a stratum is associated to a half-tree constructed in the following way:*

- (1) *Vertices correspond to horizontal cylinders. (NB: The action of the hyperelliptic involution maps each cylinder to itself and identifies pairs of saddle connections on the boundary. Gluing together these saddle connections gives a horizontally periodic hyperelliptic surface.)*
- (2) *For each pair of saddle connections interchanged by the hyperelliptic involution draw an edge between the vertices representing the cylinders they connect.*
- (3) *For all other saddle connections draw a half-edge emanating from the lone cylinder it borders.*

The total number of half-edges is $2g + |\Sigma| - 2$ where Σ is the singular set and where we count each full edge as two half-edges.

Lindsey's result is actually more general. It associates a tree to any translation surface in a hyperelliptic component; in particular the translation surface need not be horizontally periodic. In this more complicated construction, each node represents either a horizontal cylinder or a minimal component of the horizontal line flow and a new kind of half-edge is required corresponding to horizontal lines beginning at a singularity, but never terminating at a singularity.

Below are all possible half-graphs with four or fewer half-edges - i.e. the ones arising from surfaces in $\mathcal{H}(2)$ and $\mathcal{H}(1, 1)$ - and corresponding horizontally periodic surfaces.




Define the combinatorial type of a horizontally periodic translation surface in $\mathcal{H}^{hyp}(2g-2)$ or $\mathcal{H}^{hyp}(g-1, g-1)$ to be the equivalence class of translation surfaces that are related by some combination of the following three operations:

- (1) Horizontally shearing a horizontal cylinder.

- (2) Vertically dilating a horizontal cylinder.
- (3) Changing the lengths of a collection of horizontal saddle connections that is invariant under the hyperelliptic involution.

The notion of combinatorial type applies equally well to horizontal cylinders on translation surfaces.

Lemma 2.1. *Each node in the Lindsey half-tree has the combinatorial type of*



Proof. By Lindsey [10] each node corresponds to a hyperelliptic surface that is contained in one horizontal cylinder and whose hyperelliptic involution is given by rotation by π . It follows that each node has the combinatorial type shown above. \square

So that we may refer to it later the combinatorial type of the cylinder in lemma 2.1 will be referred to as hyperelliptic combinatorial type.

Theorem 2.2. *The combinatorial types of horizontally periodic translation surfaces in $\mathcal{H}^{hyp}(2g-2)$ and $\mathcal{H}^{hyp}(g-1, g-1)$ are in bijective correspondence with planar embeddings of half-trees with $2g + |\Sigma| - 2$ half-edges up to precomposition with half-tree automorphisms (here we count full edges as two half-edges).*

Proof. For the forward direction of this correspondence, take the combinatorial type of a horizontally periodic translation surface in $\mathcal{H}^{hyp}(2g-2)$ or $\mathcal{H}^{hyp}(g-1, g-1)$ and associate to it the Lindsey half-tree. We form the planar embedding (up to automorphism) by recording the cyclic ordering of half-edges attached to the node.

For the reverse direction we must take a half-tree Γ and produce a translation surface $(X_\Gamma, \omega_\Gamma)$ in a hyperelliptic component of a stratum whose associated Lindsey half-tree is Γ . Associate to each node $v \in \Gamma$ the translation surface in lemma 2.1 where n is the number of edges and half-edges attached to v . Label the edges and half-edges attached to v clockwise $\{1, \dots, n\}$. If an edge of the Lindsey tree connects the nodes v and w it will have two labels i and j coming from v and w respectively. To form $(X_\Gamma, \omega_\Gamma)$ slice open v along saddle connection i and w along saddle connection j and glue i to j . Do this for all edges in Γ .

The resulting surface has an involution given by $-I$ that fixes every horizontal cylinder. Since the surface in lemma 2.1 is hyperelliptic with hyperelliptic involution given by $-I$ the quotient of every node of the Lindsey tree is a copy of \mathbb{P}^1 . The tree structure (ignoring half-edges) of the half-tree describes how the copies of \mathbb{P}^1 glue together. Since trees are contractible it follows that $(X_\Gamma, \omega_\Gamma)/-I$ is homeomorphic to \mathbb{P}^1 . Consequently, $(X_\Gamma, \omega_\Gamma)$ is a hyperelliptic translation surface, i.e. X_Γ admits a hyperelliptic involution that acts by $-I$ on ω_Γ . To conclude that $(X_\Gamma, \omega_\Gamma)$ actually lies in $\mathcal{H}^{hyp}(2g-2)$ or $\mathcal{H}^{hyp}(g-1, g-1)$ and not in a hyperelliptic locus in another connected component it suffices to show that the surface has either one zero or two zeros that are exchanged by the hyperelliptic involution.

The translation surfaces in lemma 2.1 have a single Weierstrass point at the center of the rectangle and another $n + 1$ Weierstrass points at the midpoints of each saddle connection on the boundary (whether or not the vertex is also a Weierstrass point will depend on the parity of n). Therefore, prior to identifying saddle connections, each node v of Γ contributes $d_v + 2$ Weierstrass points (not including potential Weierstrass points at vertices) where d_v is the number of full and half-edges attached to v . When two saddle connections on nodes v and w are identified the midpoints of both saddle connections are exchanged by the hyperelliptic involution and hence cease to be fixed points. It follows that the number of Weierstrass points on $(X_\Gamma, \omega_\Gamma)$ is

$$2g + 2 = \left(\sum_{v \in V} d_v + 2 \right) - 2|E| + \#\{\text{zeros fixed by the hyperelliptic involution}\}$$

where V (resp. E) are the vertices (resp. full edges) of Γ . Since Γ is a tree $2|V| - 2|E| = 2$ and since the total number of half-edges (counting each full edge as two half edges) is $2g + |\Sigma| - 2$ we also have that $\sum_{v \in V} d_v = 2g + |\Sigma| - 2$. Therefore,

$$2g + 2 = (2g + |\Sigma| - 2 + 2|V|) - 2|E| + \#\{\text{zeros fixed by the hyperelliptic involution}\}$$

which shows that

$$|\Sigma| + \#\{\text{zeros fixed by the hyperelliptic involution}\} = 2$$

It follows that $(X_\Gamma, \omega_\Gamma)$ has either one fixed zero or two zeros that are exchanged by the hyperelliptic involution; so $(X_\Gamma, \omega_\Gamma)$ lies in $\mathcal{H}^{hyp}(2g - 2)$ or $\mathcal{H}^{hyp}(g - 1, g - 1)$. \square

Corollary 2.3. *If (X, ω) is a horizontally periodic translation surface such that every horizontal cylinder has hyperelliptic combinatorial type and whose cylinder diagram corresponds to a Lindsey tree then (X, ω) lies in a hyperelliptic component of a stratum*

3. BRANCHED COVERING CONSTRUCTIONS

A translation covering $f : (Y, \eta) \longrightarrow (X, \omega)$ is a holomorphic map $f : Y \longrightarrow X$ such that $f^*\omega = \eta$. A simple translation covering is a translation covering that is branched over the zeros of ω . The goal of this section is to develop a criterion to recognize “branched covering constructions over \mathcal{H} ” - i.e. affine invariant submanifolds all of whose elements are simple translation coverings of elements in a component \mathcal{H} of a stratum of abelian differentials.

Lemma 3.1 (Maskit-Mumford Compactness Lemma). *If $((X_n, \omega_n))_n$ is a sequence of translation surfaces in a fixed stratum that have area bounded from above and below and the length of their shortest saddle connection bounded below, then there is a convergent subsequence.*

Proof. By Maskit [11] if the length of the shortest saddle connection is bounded away from zero then the length of the shortest hyperbolic curve on X_n is bounded away from zero. By Mumford there is a convergent subsequence of X_n . Passing to this subsequence let X be the limit and let U be a precompact neighborhood of X on which the bundle of holomorphic one-forms is trivial. Since the area is bounded below and above (X_n, ω_n) eventually is contained in a bundle of compact annuli over the compact set \overline{U} . Therefore there is a

convergent subsequence. Since no saddle connection becomes short the sequence remains in the same fixed stratum. \square

Corollary 3.2. *If \mathcal{H} is a component of a stratum and $\mathcal{H}_{\geq \epsilon, A}$ is the collection of translation surface with smallest saddle connection of length at least ϵ and area bounded between $\frac{1}{A}$ and A , then these sets form a compact exhaustion of \mathcal{H} .*

Given a translation surface (X, ω) let $\pi_1(X, \omega)$ denote the fundamental group of $X - Z(\omega)$ where $Z(\omega)$ is the vanishing set of ω . A degree d simple translation covering $f : (Y, \eta) \rightarrow (X, \omega)$ defines a monodromy representation, i.e. a representation $\rho : \pi_1(X, \omega) \rightarrow S_d$ where S_d is the symmetric group on d elements and the representation is well-defined up to post-composition by an inner automorphism of S_d . Conversely given a representation $\rho : \pi_1(X, \omega) \rightarrow S_d$ it is possible to define a Riemann surface Y' such that $Y' \rightarrow X - Z(\omega)$ is a cover of prescribed monodromy. The Riemann extension theorem implies that Y' may be completed to a compact Riemann surface and the covering map $\pi : Y' \rightarrow X - Z(\omega)$ may be extended to a holomorphic branched covering map $\pi : Y \rightarrow X$. The map π is well-defined up to pre-composition with an element of the deck group. If g is an element of the deck group then $\pi \circ g = \pi$ and $\pi^* \omega = g^* \pi^* \omega$. Therefore, $(Y, \eta := \pi^* \omega)$ is well-defined given (X, ω) and the monodromy representation. We conclude that a degree d simple translation covering $f : (Y, \eta) \rightarrow (X, \omega)$ and a monodromy representation $\rho : \pi_1(X, \omega) \rightarrow S_d$ up to conjugacy define the same datum.

Suppose that \mathcal{H} contains genus g surfaces with n zeros. Let $\text{Mod}_{g,n}$ be the mapping class group of a smooth genus g Riemann surface with n punctures (elements of this group may permute the punctures). On the Teichmüller space $\text{Teich}_{g,n}$ each puncture is labelled. Suppose that \mathcal{H} is a component of $\mathcal{H}(\alpha) := \mathcal{H}(\alpha_1^{n_1}, \dots, \alpha_k^{n_k})$ where α_i denotes the degree of a zero and n_i specifies the number of zeros of that degree. Let $\text{Teich}_{g,\alpha}$ be the collection of markings on smooth genus g surfaces with n punctures, but instead of labelling each puncture separately the punctures are labelled by symbols $\{\alpha_i, \dots, \alpha_k\}$ with n_i punctures having label α_i . Let $\text{Mod}(\mathcal{H})$ be the subgroup of $\text{Mod}_{g,n}$ that stabilizes the collection of punctures of label α_i for each i . $\text{Mod}(\mathcal{H})$ is a finite index subgroup in $\text{Mod}_{g,n}$.

Fix a reference point $(S, \alpha) \in \mathcal{H}$ and let $\rho : (S, \alpha) \rightarrow S_d$ be a monodromy representation. Let $\tilde{\mathcal{H}}$ be the cover of \mathcal{H} whose elements are translation surfaces in \mathcal{H} with a marking. To be explicit, the elements of $\tilde{\mathcal{H}}$ are isotopy classes of quasiconformal maps $\phi : (S, \alpha) \rightarrow (X, \omega)$ such that zeros of α are sent to zeros of ω of the same degree. The quotient of $\tilde{\mathcal{H}}$ by $\text{Mod}(\mathcal{H})$ coincides with \mathcal{H} . Define $\text{Mod}(\mathcal{H}, \rho)$ to be the collection of elements g in $\text{Mod}(\mathcal{H})$ such that $\rho \circ g_*$ is identical to ρ up to post-composition by an inner automorphism of S_d . Since the collection of monodromy representations of $\pi_1(S, \alpha)$ is finite it follows that $\text{Mod}(\mathcal{H}, \rho)$ has finite index in $\text{Mod}(\mathcal{H})$. Let $\tilde{\tilde{\mathcal{H}}}$ be the finite cover of \mathcal{H} given by $\tilde{\mathcal{H}}/\text{Mod}(\mathcal{H}, \rho)$, i.e. the collection of $(X, \omega) \in \mathcal{H}$ together with an equivalence class of markings so that two markings $\phi : (S, \alpha) \rightarrow (X, \omega)$ and $\psi : (S, \alpha) \rightarrow (X, \omega)$ are equivalent if $\rho \circ \phi_*^{-1}$ and $\rho \circ \psi_*^{-1}$ define the same monodromy representation of $\pi_1(X, \omega)$ up to post-composition by an inner automorphism of S_d .

Let (Y, η) be the branched cover associated to the monodromy representation $\rho : \pi_1(S, \alpha) \rightarrow S_d$ and suppose that it is connected. Suppose that \mathcal{H}' is the component of the stratum of

abelian differentials that contains (Y, η) . Let $F : \tilde{\mathcal{H}} \rightarrow \mathcal{H}'$ be the map that associates a marking $\phi : (S, \alpha) \rightarrow (X, \omega)$ to the branched cover (Z, ζ) corresponding to the monodromy representation $\rho \circ \phi_*^{-1} : \pi_1(X, \omega) \rightarrow S_d$. It is clear that F is continuous by working in period coordinates and it is clear that F factors through $\tilde{\mathcal{H}}$. The following commutative diagram summarizes the situation

$$\begin{array}{ccc} & & \mathcal{H}' \\ & \nearrow F & \\ \tilde{\mathcal{H}} & \rightarrow & \tilde{\mathcal{H}} \\ & \searrow & \\ & & \mathcal{H} \end{array}$$

The following succinct proof was suggested by Alex Wright

Lemma 3.3. *$F : \tilde{\mathcal{H}} \rightarrow \mathcal{H}'$ is a closed, equivariant map; therefore $F(\tilde{\mathcal{H}})$ is an affine invariant submanifold of \mathcal{H}' .*

Proof. The $\mathrm{GL}_2 \mathbb{R}$ action on translation surfaces induces a $\mathrm{GL}_2 \mathbb{R}$ action on marked translation surfaces by post-composing the marking with the element of $\mathrm{GL}_2 \mathbb{R}$. Suppose that $f : (Y, \eta) \rightarrow (X, \omega)$ is a simple translation covering with monodromy given by $\rho \circ \phi_*^{-1}$ for some marking $\phi : (S, \alpha) \rightarrow (X, \omega)$. To show that F is equivariant it suffices to show for any $g \in \mathrm{GL}_2 \mathbb{R}$ that $F(g \cdot (\phi : (S, \alpha) \rightarrow (X, \omega))) = g \cdot (Y, \eta)$. Notice that $gfg^{-1} : g(Y, \eta) \rightarrow g(X, \omega)$ is a holomorphic map with monodromy given by $\rho \circ \phi_*^{-1} \circ g_*^{-1}$, which shows equivariance. To show that F is closed it suffices to show that $F : \tilde{\mathcal{H}} \rightarrow \mathcal{H}'$ is proper. This is immediate since $F^{-1}(\mathcal{H}'_{\epsilon, A}) \subseteq \tilde{\mathcal{H}}_{\frac{\epsilon}{d}, dA}$. \square

Corollary 3.4 (Definition of (Simple) Branched Covering Construction). *If \mathcal{M} is a connected affine invariant submanifold containing an \mathcal{M} -generic translation surface (Y, η) that is a simple translation covering of $(X, \omega) \in \mathcal{H}$, then \mathcal{M} is contained in $F(\tilde{\mathcal{H}})$ where F is defined using the monodromy representation of the simple translation covering. We say that \mathcal{M} is a (simple) branched covering construction of \mathcal{H} .*

Since the only branched covering constructions we consider are simple ones we will suppress the word “simple” throughout the remainder of the text.

4. CYLINDER DEFORMATIONS

Throughout this section \mathcal{M} will be an affine invariant submanifold in a component \mathcal{H} of a stratum. Suppose that (X, ω) is a translation surface in an affine invariant submanifold \mathcal{M} . Let C and C' be two cylinders on (X, ω) with core curves γ_C and $\gamma_{C'}$. If γ_C and $\gamma_{C'}$ are parallel in some neighborhood $U \subseteq \mathcal{M}$ of (X, ω) , then C and C' are said to be \mathcal{M} -equivalent.

Theorem 4.1 (Cylinder Proportion Theorem, Proposition 3.2, Nguyen-Wright [19]). *If C and C' are two \mathcal{M} -equivalent cylinders and \mathcal{V} is any equivalence class of cylinders, then*

$$\frac{|C \cap \mathcal{V}|}{|C|} = \frac{|C' \cap \mathcal{V}|}{|C'|}$$

where $|\cdot|$ denotes area.

Applying the matrix $u_t := \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}$ to a horizontal \mathcal{M} -equivalence class of horizontal cylinders \mathcal{C} will be called (horizontally) shearing \mathcal{C} . Applying the matrix $a_t := \begin{pmatrix} 1 & 0 \\ 0 & e^t \end{pmatrix}$ will be called (vertically) dilating \mathcal{C} .

Theorem 4.2 (Cylinder Deformation Theorem, Wright [22]). *Let $(X, \omega) \in \mathcal{M}$ be a translation surface and let \mathcal{C} be an \mathcal{M} -equivalence class of horizontal cylinders on (X, ω) . Horizontally shearing and vertically dilating \mathcal{C} by t remains in \mathcal{M} for all t .*

A special feature of the hyperelliptic component of a stratum is that if two horizontal cylinders share a horizontal saddle connection then they share exactly two - one on the top of each cylinder and one on the bottom of each cylinder. These saddle connections are exchanged by the hyperelliptic involution.

Lemma 4.3 (Standard Position). *Suppose that (X, ω) is a translation surface in a hyperelliptic component of a stratum and suppose that C and D are adjacent horizontal cylinders that belong to distinct equivalence classes \mathcal{C}_1 and \mathcal{C}_2 respectively. It is possible to shear \mathcal{C}_1 and \mathcal{C}_2 so that there is a vertical cylinder V that simply intersects C and D , intersects no other horizontal cylinders, and that contains the two horizontal saddle connections on the boundary of C and D . Given two such cylinders this position will be called standard position. Alternatively, it is possible to only shear \mathcal{C}_2 so that there is a cylinder T transverse to the horizontal that simply intersects C and D , intersects no other horizontal cylinders, and contains the two horizontal saddle connections on the boundary of C and D . To move to this configuration we will say that we have put C and D in transverse standard position while fixing C .*

Proof. Let s_1 and s_2 be the horizontal saddle connections lying on the boundary of C and D . By Theorem 4.2 shear \mathcal{C}_1 so that s_1 lies directly over s_2 ; then shear \mathcal{C}_2 so that s_2 lies directly over s_1 . Recall that s_1 and s_2 are exchanged by the hyperelliptic involution and hence have identical lengths. Choose V to be the vertical cylinder passing through s_1 and s_2 . \square

If C is a cylinder call the distance h_C from one boundary of the cylinder to the other its height. Let γ_C^* be the cohomology class that is dual to the core curve of C under the intersection pairing. This cohomology class requires specifying an orientation on γ_C . Usually this orientation will not matter, but we will establish the conventions that when C is horizontal the orientation is left to right, when C is vertical it is top to bottom, and when C_1, \dots, C_n are all \mathcal{M} -equivalent cylinders we will assume that the holonomy vectors of the core curves point in the same directions.

Let $\mathcal{C}_1, \dots, \mathcal{C}_n$ be an enumeration of the \mathcal{M} -equivalence classes of horizontal cylinders on (X, ω) . For each equivalence class \mathcal{C}_i there is an element of the tangent space called the standard shear which is defined to be $u_{\mathcal{C}_i} = \sum_{c \in \mathcal{C}_i} h_c \gamma_c^*$. A reformulation of the cylinder deformation theorem is that the standard shear is always in the tangent space of \mathcal{M} at (X, ω) .

Let \mathcal{C} denote the collection of all horizontal cylinders on (X, ω) . The twist space of \mathcal{M} at (X, ω) is defined to be

$$\text{Twist}_{(X, \omega)} \mathcal{M} := \text{span}_{\mathbb{R}} (\gamma_c^*)_{c \in \mathcal{C}} \cap T_{(X, \omega)}^{\mathbb{R}} \mathcal{M}$$

where $T_{(X, \omega)}^{\mathbb{R}} \mathcal{M} = T_{(X, \omega)} \mathcal{M} \cap H^1(X, Z(\omega); \mathbb{R})$ where $T_{(X, \omega)} \mathcal{M}$ has been identified with a subspace of $H^1(X, Z(\omega); \mathbb{C})$. The standard shears are always in the twist space. Define the cylinder preserving space, denoted $\text{Pres}_{(X, \omega)} \mathcal{M}$, to be all elements of $T_{(X, \omega)}^{\mathbb{R}} \mathcal{M}$ that pair trivially with every element of $(\gamma_c)_{c \in \mathcal{C}}$ under the intersection pairing. It is clear that $\text{Twist}_{(X, \omega)} \mathcal{M} \subseteq \text{Pres}_{(X, \omega)} \mathcal{M}$. The following theorem establishes that there is always a translation surface (X, ω) in an affine invariant submanifold \mathcal{M} where $\text{Twist}_{(X, \omega)} \mathcal{M} = \text{Pres}_{(X, \omega)} \mathcal{M}$.

Theorem 4.4 (Lemma 8.6, Wright [22]). *Twist_(X, ω) M = Pres_(X, ω) M whenever (X, ω) has the maximum number of horizontal cylinders in M.*

The next theorem indicates why having $\text{Twist}_{(X, \omega)} \mathcal{M} = \text{Pres}_{(X, \omega)} \mathcal{M}$ is special. In particular, it says that whenever equality is achieved (X, ω) has at least $\text{rk}(\mathcal{M})$ many \mathcal{M} -equivalence classes and the twist space projects to a Lagrangian in $p(T_{(X, \omega)} \mathcal{M})$.

Theorem 4.5 (Lemma 8.12, Wright [22]). *If (X, ω) is a translation surface in M and Twist_(X, ω) M = Pres_(X, ω) M then span_ℝ p(u_{c_i})_{i=1}ⁿ is a Lagrangian subspace of p(T_(X, ω)^ℝ M) where {C₁, ..., C_n} is an enumeration of the M-equivalence classes of horizontal cylinders and u_{c_i} is the standard shear. In particular, (X, ω) contains at least rk(M) distinct M-equivalence classes of horizontal cylinders.*

The combination of the previous two theorems is an engine that allows us to convert the rank of an affine invariant submanifold into geometric information that picks out a translation surface where a large dimensional subspace, the twist space, of the tangent space is geometrically meaningful.

Theorem 4.6 (Twist Space Decomposition Theorem, Theorem 4.7, Mirzakhani-Wright [17]).

Let (X, ω) be a translation surface in an affine invariant submanifold M. Let C₁, ..., C_d be an enumeration of the M-equivalence classes of horizontal cylinders.

- (1) *If v ∈ Twist_(X, ω) M then v = ∑_{i=1}^d v_i where v_i ∈ Twist_(X, ω) M ∩ span_ℝ (γ_c^{*})_{c ∈ C_i}.*
- (2) *If v_i ∈ Twist_(X, ω) M ∩ span_ℝ (γ_c^{*})_{c ∈ C_i} then v_i ∈ ℝ · u_{c_i} ⊕ ker p where u_{c_i} is the standard shear.*

The last result regarding cylinder deformations that we need is the statement that given a collection of d \mathcal{M} -equivalence classes of cylinders, where d is no bigger than the rank of \mathcal{M} , it is possible to perturb the translation surface so that one \mathcal{M} -equivalence class becomes disjoint and vertical and all others remain horizontal. This result is crucial establishing that all higher rank affine invariant submanifolds of complex dimension four in hyperelliptic components of strata are branched covering constructions over $\mathcal{H}(2)$.

Theorem 4.7 (Perturbation Theorem). *Suppose that (X, ω) is a translation surface in an affine invariant submanifold M. Suppose that C₁, ..., C_d are M-equivalence classes of horizontal cylinders such that {p(u_{C₁}), ..., p(u_{C_d})} spans a d dimensional subspace. Define C := C₁ ∪ ... ∪ C_{d-1}. There is a piecewise smooth path f : [0, 1] → M such that f(0) = (X, ω) and along the path*

- (1) All cylinders in \mathcal{C} persist and are horizontal.
- (2) The cylinders in \mathcal{C}_d persist, become nonhorizontal for $t > 0$, and vertical at $t = 1$.
- (3) At all points along the path any cylinder \mathcal{M} -equivalent to \mathcal{C}_d is disjoint from any point \mathcal{M} -equivalent to \mathcal{C}_i for all $i \in \{1, \dots, d-1\}$.

Proof. Let γ_i be the core curve of some cylinder in \mathcal{C}_i for each $i \in \{1, \dots, d\}$. Consider the linear functionals $(\gamma_i)_{i=1}^{d-1}$ on $T_{(X,\omega)}\mathcal{M}$. Since the linear functionals factor through $p : T_{(X,\omega)}\mathcal{M} \rightarrow H^1(X; \mathbb{C})$ we see that the intersection of the kernel of these functionals on $p(T_{(X,\omega)}\mathcal{M})$ is at least dimension $(2r-d)+1$. Therefore there is a vector $v \in T_{(X,\omega)}\mathcal{M}$ that is not in the kernel of p , not in the cylinder preserving space, and such that $v(\gamma_i) = 0$ for $1 \leq i \leq d-1$.

Consider the path $(X, \omega) + tv$ for $t \geq 0$. This path is well-defined and remains in \mathcal{M} for some range $t \in [0, T]$. Since $v(\gamma_i) = 0$ for each $1 \leq i \leq d-1$ it follows by definition of \mathcal{M} -equivalence that \mathcal{C}_i persist (perhaps after decreasing T) and remain horizontal for $1 \leq i \leq d-1$. Since v is not in the cylinder preserving space we see that, perhaps after decreasing T , \mathcal{C}_d also persists and becomes non-horizontal. By perhaps decreasing T again we may suppose by Mirzakhani-Wright Lemma 5.5 [17] that any cylinder \mathcal{M} -equivalent to \mathcal{C}_d remains disjoint from any cylinder \mathcal{M} -equivalent to \mathcal{C}_i at all points along the path.

Now horizontally shear $(X, \omega) + Tv$ until \mathcal{C}_d becomes vertical. The equations on period coordinates cutting out \mathcal{M} may be parallel translated along this path and so we see that along this path no new cylinders become \mathcal{M} -equivalent to cylinders in \mathcal{C}_i for any i . \square

Lemma 4.8. *Let \mathcal{M} be an affine invariant submanifold. Suppose that for any horizontally periodic $(X, \omega) \in \mathcal{M}$ such that $\text{Twist}_{(X,\omega)}\mathcal{M} = \text{Pres}_{(X,\omega)}\mathcal{M}$, the heights of any two \mathcal{M} -equivalent horizontal cylinders are identical. Then the heights of any two equivalent cylinders on any translation surface in \mathcal{M} are identical.*

Proof. Let \mathcal{C} be an equivalence class of horizontal cylinders on $(X_0, \omega_0) \in \mathcal{M}$. Let $\epsilon > 0$. Consider the following iterative process:

- (1) If (X_i, ω_i) is horizontally periodic and $\text{Twist}_{(X_i,\omega_i)} = \text{Pres}_{(X_i,\omega_i)}$ then end the process. Otherwise, Smillie-Weiss [20] states that the horocycle flow accumulates on a horizontally periodic translation surface (Y_i, η_i) , on which the cylinders in \mathcal{C} persist, remain \mathcal{M} -equivalent, and so that each cylinder has a height $\frac{\epsilon}{2 \cdot (g+|\Sigma|-1)}$ -close to its height on (X_i, ω_i) .
- (2) If $\text{Twist}_{(Y_i,\eta_i)} = \text{Pres}_{(Y_i,\eta_i)}$ then end the process. Otherwise there is an element v in the cylinder preserving space that fails to be in the twist space. Flowing in the $\sqrt{-1} \cdot v$ direction for an arbitrarily small positive time leads to a surface (X_{i+1}, ω_{i+1}) on which the cylinders in \mathcal{C} persist, are \mathcal{M} -equivalent and, have heights that are $\frac{\epsilon}{2 \cdot (g+|\Sigma|-1)}$ -close to their heights on (Y_i, η_i) ; but where the horizontal cylinders from (Y_i, η_i) although they persist, do not cover (X_{i+1}, ω_{i+1}) . Now return to step 1.

Since the number of cylinders increases with each iteration and the largest possible number of horizontal cylinders is $g+|\Sigma|-1$ the process terminates after at most $g+|\Sigma|-1$ cycles. At the end of the process, \mathcal{C} is a collection of \mathcal{M} -equivalent cylinders on a translation surface (X, ω) with heights ϵ -close to their original heights and where $\text{Twist}_{(X,\omega)}\mathcal{M} = \text{Pres}_{(X,\omega)}\mathcal{M}$.

By hypothesis, these cylinders on (X, ω) have identical heights. Therefore, the cylinders in \mathcal{C} on (X_0, ω_0) all have heights that are ϵ -close to one another for arbitrary ϵ . \square

5. ODD DIMENSIONAL ORBIT CLOSURES

Throughout this section \mathcal{M} will be an affine invariant submanifold in $\mathcal{H}^{hyp}(g-1, g-1)$ of odd complex dimension $2r+1$. Let $\eta_{(X, \omega)}$ denote the kernel of the map $p : T_{(X, \omega)}\mathcal{M} \rightarrow H^1(X, \mathbb{C})$. We will identify η uniquely (up to scaling) with an element of $T_{(X, \omega)}^{\mathbb{R}}\mathcal{M}$, which is defined to be tangent vectors that correspond to real cohomology classes.

Theorem 5.1. *Let (X, ω) be a translation surface in \mathcal{M} with at least one horizontal cylinder. The following are equivalent:*

- (1) (X, ω) has $g+1$ horizontal cylinders (equivalently, the Lindsey half-tree is a tree and not just a half-tree).
- (2) $\text{Twist}_{(X, \omega)}\mathcal{M} = \text{Pres}_{(X, \omega)}\mathcal{M}$
- (3) The relative deformation η is contained in the twist space.

When any of these equivalent conditions holds label the cylinders $\{c_0, \dots, c_g\}$ and the core curves $\{\gamma_0, \dots, \gamma_g\}$, it follows that $\eta = \sum_{i=0}^g (-1)^{d(c_0, c_i)} \gamma_i^*$ where $d(c_i, c_0)$ is the distance between c_i and c_0 in the Lindsey tree.

Proof. (1 \Rightarrow 2) By Theorem 4.4 if (X, ω) has $g+1$ horizontal cylinders then $\text{Twist}_{(X, \omega)}\mathcal{M} = \text{Pres}_{(X, \omega)}\mathcal{M}$.

(2 \Rightarrow 3) Since any relative deformation fixes the core curves of every cylinder it follows that if $\text{Twist}_{(X, \omega)}\mathcal{M} = \text{Pres}_{(X, \omega)}\mathcal{M}$ then η is contained in $\text{Twist}_{(X, \omega)}\mathcal{M}$.

(3 \Rightarrow 1) Now suppose that η is contained in the twist space, i.e $\eta = \sum_c a_c \gamma_c^*$ where the sum is taken over the collection of horizontal cylinders and $a_c \in \mathbb{R}$. Since (X, ω) is a translation surface in a hyperelliptic connected component whenever v and v' are adjacent cylinders there is an absolute period contained in $v \cup v'$ that intersects γ_v and $\gamma_{v'}$ exactly once with the same orientation. Since this period must be unchanged by the relative deformation η it follows that $a_v + a_{v'} = 0$ for any two adjacent cylinders v and v' . Since the vertices are arranged in a tree we have that up to scaling the purely relative deformation is $\sum_c (-1)^{d(c, c_0)} \gamma_c^*$ where c_0 is some fixed cylinder. Finally, the Lindsey tree of (X, ω) cannot have any half-edges since they yield nonzero elements of absolute homology that are supported in a single cylinder and hence will be altered by $\sum_c (-1)^{d(c, c_0)} \gamma_c^*$. Therefore, the Lindsey tree of (X, ω) is a tree (not just a half-tree) and (X, ω) has $g+1$ horizontal cylinders. \square

Theorem 5.2. *If \mathcal{M} has rank $r > 1$ and if $(X, \omega) \in \mathcal{M}$ has $g+1$ horizontal cylinders then (X, ω) has $r+1$ equivalence classes. If $\mathcal{C}_0, \dots, \mathcal{C}_r$ is an enumeration of the equivalence classes then*

$$\text{Twist}_{(X, \omega)}\mathcal{M} = \text{span}_{\mathbb{R}}\{u_{\mathcal{C}_0}, \dots, u_{\mathcal{C}_r}\}$$

where u_{C_i} is the standard shear of C_i . Any two cylinders in the same equivalence class have identical heights and are an even distance apart in the Lindsey tree. Moreover, \mathcal{M} is defined over \mathbb{Q} .

Proof. If (X, ω) has $g + 1$ cylinders then $\text{Twist}_{(X, \omega)} \mathcal{M} = \text{Pres}_{(X, \omega)} \mathcal{M}$. Since \mathcal{M} is higher rank it follows that there are at least two \mathcal{M} -equivalence classes of horizontal cylinders and so Theorem 5.1 implies that η is not supported on a single \mathcal{M} -equivalence class. By the twist space decomposition theorem (Theorem 4.6) it follows that the only element of the twist space supported on a single equivalence class is the standard shear and hence η is a real linear combination of standard shears, i.e.

$$\sum_{i=0}^g (-1)^{d(c_i, c_0)} \gamma_i^* = \sum_{i=1}^n a_i \sum_{c \in C_i} h_c \gamma_c^*$$

Since γ_c^* are all linearly independent in $T_{(X, \omega)} \mathcal{H}$ where \mathcal{H} is the component of the stratum containing (X, ω) it follows that any two equivalent cylinders have the same height and are an even distance apart in the Lindsey tree. Finally we see that $\text{Twist}_{(X, \omega)} \mathcal{M}$ is spanned by standard shears and its projection to absolute cohomology has a one-dimensional kernel. It follows that (X, ω) has $r + 1$ equivalence classes of cylinders

By theorem 7.1 [22] to show that \mathcal{M} is defined over \mathbb{Q} it suffices to show that the ratio of lengths of core curves of any two equivalent horizontal cylinders is always rational. Notice that the only element of the twist space supported on a single equivalence class is the standard shear. This implies that the ratio of moduli of any two equivalent cylinders is rational. Since the heights of any two equivalent cylinders are identical the result follows. \square

Corollary 5.3. *Any two equivalent horizontal cylinders on any translation surface in \mathcal{M} have identical heights when \mathcal{M} is higher rank.*

Proof. This is immediate from Lemma 4.8 and Theorem 5.2. \square

Corollary 5.4. *If $(X, \omega) \in \mathcal{M}$ is a translation surface with at least one horizontal cylinder and \mathcal{M} is higher rank then the twist space is spanned by standard shears of the \mathcal{M} -equivalence classes of horizontal cylinders.*

Proof. By Theorem 5.1, Theorem 5.2 establishes this result when the twist space contains η . When the twist space does not contain η the result is immediate by the twist space decomposition theorem. \square

Theorem 5.5. *Suppose that \mathcal{M} has rank one and (X, ω) has $g + 1$ cylinders. Suppose furthermore that C_0, C_1 is a partition of the cylinders so that $T_{(X, \omega)}^{\mathbb{R}} \mathcal{M}$ contains $\sum_{c \in C_0} h_c \gamma_c^*$.*

Then any two cylinders in C_0 (resp. C_1) have identical heights and are an even distance apart in the Lindsey tree. Moreover, \mathcal{M} is defined over \mathbb{Q} and hence is a branched covering construction of $\mathcal{H}(0, 0)$.

Proof. The proof is identical to that of Theorem 5.2. \square

6. A PARTIAL COMPACTIFICATION OF STRATA OF ABELIAN DIFFERENTIALS

A natural partial compactification of a stratum of abelian differentials is the bundle of stable finite volume abelian differentials over the Deligne-Mumford compactification of moduli space. However, it is often more natural from the perspective of flat geometry to consider a quotient of this space that ignores components of the underlying curve on which the stable one-form vanishes. For the remainder of the paper this quotient will be called the partial compactification of a stratum. It was introduced in McMullen [16] and extensively studied in Mirzakhani-Wright [17].

The following example of convergence to the boundary is Example 3.1 of [17]. Suppose that \mathcal{M} is an affine invariant submanifold and let (X, ω) be a translation surface in \mathcal{M} with an \mathcal{M} -equivalence class of horizontal cylinders \mathcal{C} . Suppose that \mathcal{C} does not cover (X, ω) and that the union of cylinders in \mathcal{C} contains a vertical saddle connection. Let (X_t, ω_t) be (X, ω) with \mathcal{C} vertically shrunk by e^t . By the cylinder deformation theorem it follows that (X_t, ω_t) is a smooth path in \mathcal{M} . This sequence converges to a translation surface $(X_\infty, \omega_\infty)$ on the boundary of \mathcal{M} . To form $(X_\infty, \omega_\infty)$ use the following procedure:

- (1) Delete every cylinder in \mathcal{C} from (X, ω) to form a translation surface with boundary. The boundary is a collection of saddle connections.
- (2) If there is a point p on the boundary of $(X, \omega) - \mathcal{C}$ that is joined to a zero or marked point of (X, ω) by a vertical line that is completely contained in \mathcal{C} then mark p . On the boundary translation surface these points will be either marked points or zeros of the boundary holomorphic one-form. Adding in these marked points may divide saddle connections into several smaller saddle connections.
- (3) If two saddle connections on the boundary of $(X, \omega) - \mathcal{C}$ were connected by a vertical line that was completely contained in \mathcal{C} then glue the two saddle connections together. The resulting translation surface is $(X_\infty, \omega_\infty)$.

This construction is called a horizontal cylinder collapse. The analogous construction with an \mathcal{M} -equivalence class of vertical cylinders will be called a vertical cylinder collapse.

Theorem 6.1 (Mirzakhani-Wright [17]; Proposition 2.3). *Given a sequence $(X_n, \omega_n, \Sigma_n)$ of translation surfaces converging to (X, ω, Σ) there are collapse maps $f_n : X_n \rightarrow X$ such that*

- (1) *There is a neighborhood U_n of Σ_n so that $f_n : X_n - U_n \rightarrow X$ is a diffeomorphism onto its image with inverse g_n .*
- (2) *The injectivity radius of U_n goes to zero uniformly in n .*
- (3) *$g_n^* \omega_n$ converges to ω in the compact open topology.*

Define the space of vanishing cycles to be

$$V_n = \ker (f_n : H_1(X_n, \Sigma_n; \mathbb{C}) \rightarrow H_1(X, \Sigma; \mathbb{C}))$$

For large enough n this space is constant and will be called V .

Returning to the example of the horizontal cylinder collapse: the space of vanishing cycles will be the subspace spanned by the heights of the horizontal cylinders in \mathcal{C} and by any homology classes that have a representative supported in a subsurface that collapses to a point in the limit.

Theorem 6.2 (Degeneration Theorem; Theorem 2.7, Mirzakhani-Wright [17]). *Let \mathcal{M} be an affine invariant submanifold. Let $(X_n, \omega_n, \Sigma_n)$ be translation surfaces in \mathcal{M} converging to $(X_\infty, \omega_\infty, \Sigma)$. Let (Y, η) be a component of $(X_\infty, \omega_\infty)$ and let $\iota : (Y, \eta) \hookrightarrow (X_\infty, \omega_\infty)$ be the inclusion map. Let V be the space of vanishing cycles. The $\mathrm{GL}_2 \mathbb{R}$ orbit closure of (Y, η) is an affine invariant submanifold \mathcal{M}' whose tangent space is*

$$T_{(Y, \eta)} \mathcal{M}' = \iota^* (T_{(X_n, \omega_n)} \cap \mathrm{Ann}(V_n))$$

where $T_{(X_n, \omega_n)} \cap \mathrm{Ann}(V_n)$ has been identified with the tangent space at the boundary by parallel transport. As a consequence, $\dim_{\mathbb{C}} \mathcal{M}' < \dim_{\mathbb{C}} \mathcal{M}$ and $\mathrm{rk}(\mathcal{N}) \leq \mathrm{rk}(\mathcal{M})$ where the inequality is strict if $\mathrm{rel}(\mathcal{M}) = 0$

Throughout this paper the only degenerations that we will use are horizontal and vertical cylinder collapses.

Corollary 6.3 (Mirzakhani-Wright for Cylinder Collapses). *Let \mathcal{M} be an affine invariant submanifold and let $(X, \omega) \in \mathcal{M}$ be a translation surface with an \mathcal{M} -equivalence class of horizontal cylinders \mathcal{C} . Let (X_t, ω_t) be the horizontal cylinder collapse. Let (Y, η) be a component of the boundary translation surface and let \mathcal{N} be its orbit closure. Let W be the \mathbb{R} -linear span of the core curves of horizontal cylinders on $(X_\infty, \omega_\infty)$. By abuse of notation identify W with a subspace of $H_1(X, \mathbb{R})$ by identifying the cylinder core curves on (Y, η) with their unique pre-images on (X, ω) . The twist space of $(X_\infty, \omega_\infty)$ is isomorphic to $\mathrm{Twist}_{(X, \omega)} \mathcal{M} \cap W^*$ where W^* is the dual of W under the intersection pairing.*

Proof. The twist space is the intersection of the tangent space with the subspace of cohomology spanned by the duals of the core curves of horizontal cylinders. Since the collection of horizontal cylinders is invariant along a horizontal cylinder collapse the twist space is invariant along a horizontal cylinder collapse. The result now follows from the degeneration theorem. \square

7. DEGENERATING TO THE BOUNDARY IN HYPERELLIPTIC COMPONENTS OF STRATA

Throughout this section we will assume that (X, ω) is a horizontally periodic translation surface in an affine invariant submanifold \mathcal{M} that is contained in a hyperelliptic component of a stratum. Let Γ be the Lindsey tree of (X, ω) and let J denote the hyperelliptic involution. Suppose moreover that \mathcal{C} is an \mathcal{M} -equivalence class of cylinders that does not contain (X, ω) in its union. We will study the boundary translation surfaces when \mathcal{C} is collapsed.

Lemma 7.1 (Vertical Cylinder Collapse Lemma). *If \mathcal{C} is a collection of vertical cylinders and contains a horizontal saddle connection, then collapsing it passes to a disjoint union of translation surfaces in hyperelliptic components of strata.*

Proof. By Corollary 2.3 a translation surface belongs to a hyperelliptic component if and only if it is constructed in the following way:

- (1) For each node in the Lindsey tree Γ of degree n associate a translation surface in a hyperelliptic component that is contained in a single horizontal cylinder; this translation surface has the combinatorial type of the following surface.

Rephrased, it suffices to show the following. Suppose that C is a single cylinder of hyperelliptic type. Let s_1, \dots, s_n be saddle connections on the boundary of C . Let G be a graph with n vertices labelled $\{1, \dots, n\}$. Connect vertices i and j if s_i and $J(s_j)$ are connected by a vertical line. Then G is a disjoint union of trees. This follows immediately from the following lemma.

Lemma. *Suppose that there is a graph G with vertices labelled $\{1, \dots, n\}$. Let C_i be the set of vertices connected to vertex i . Suppose that for all i there is an increasing subset of $\{1, \dots, n\}$ (perhaps wrapping around 0) that we will denote $I_i = (k_i, k_i + 1, \dots, \ell_i)$ such that*

- (1) $C_i \subseteq I_i$ for all i .
- (2) $I_i \cap I_{i+1} = \{k_i\} = \{\ell_{i+1}\}$ if $n > 2$.

Then G is a disjoint union of trees.

Proof. Proceed by induction on n . The $n = 2$ base case is trivial. Now suppose that $n > 2$. Suppose to a contradiction that G contains a cycle. Let γ be the shortest cycle in G . If the cycle fails to contain every vertex, then delete the vertices not contained in γ from G . The induction hypothesis implies that the resulting graph cannot contain a cycle, which is a contradiction. Suppose then without loss of generality that γ involves every vertex.

If the degree of a vertex i is greater than two then we may suppose that $C_i = \{k_i, k_{i+1}, \dots, \ell_i\}$. By the hypotheses, vertices $k_i + 1, \dots, \ell_i - 1$ only connect to vertex i . Since γ is the shortest cycle in G it does not pass through vertices $k_i + 1, \dots, \ell_i - 1$ contrary to our assumption that γ passes through every vertex. It follows that every vertex in G has degree two.

Since every vertex in G has degree two and appears exactly once in γ it follows that $C_i = \{k_i, k_i + 1\}$ and $C_{i+1} = \{k_i - 1, k_i\}$ for all i . Therefore, the path γ is $(\gamma_1, \gamma_2, \dots, \gamma_n) = (1, k_1, 2, k_1 - 1, 3, k_1 - 2, \dots)$. Notice that the order of the odd vertices is $(\gamma_1, \gamma_3, \dots) = (1, 2, \dots, n)$. Since every vertex appears exactly once in γ this implies that $n = 2m + 1$ and the path γ is $(1, m + 2, 2, m + 3, \dots, m, 2m + 1, m + 1)$. However the order of the even vertices must be $(\gamma_2, \gamma_4, \dots) = (m + 2, m + 1, m, \dots)$. Therefore the cyclic order $(1, 2, \dots, n)$ and the cyclic order $(n, n - 1, \dots, 1)$ must be the same order. This only occurs when $n = 2$. But we have supposed that $n > 2$, which is a contradiction. \square

\square

Lemma 7.3 (Twist Space Degeneration Lemma). *Suppose that (X, ω) is a horizontally periodic translation surface in a higher rank affine invariant submanifold \mathcal{M} in a hyperelliptic component of a stratum. Let \mathcal{C} be an \mathcal{M} -equivalence class of cylinders as in either the horizontal or vertical cylinder collapse lemma. Collapse \mathcal{C} and let $(X_\infty, \omega_\infty)$ be the boundary translation surface. Let (Y, η) be a component of $(X_\infty, \omega_\infty)$ and suppose that its orbit closure is \mathcal{N} . Let k be the number of pairwise \mathcal{M} -inequivalent horizontal cylinders that persist on (Y, η) .*

- (1) $\dim_{\mathbb{R}} \text{Twist}_{(Y, \eta)} \mathcal{N} = k$.
- (2) *If \mathcal{N} is higher rank then two horizontal cylinders on (Y, η) are \mathcal{N} -equivalent if and only if the corresponding cylinders on (X, ω) are \mathcal{M} -equivalent.*
- (3) *If \mathcal{N} is rank one and the twist space is two dimensional then \mathcal{N} is a branched covering construction of $\mathcal{H}(0, 0)$. If C and D are \mathcal{M} -equivalent horizontal cylinders*

and both persist on (Y, η) then C and D have identical heights and are not adjacent on (X, ω) .

- (4) If $\text{Twist}_{(X, \omega)} \mathcal{M} = \text{Pres}_{(X, \omega)} \mathcal{M}$ and a horizontal cylinder from every \mathcal{M} -equivalence class persists on (Y, η) then $\dim_{\mathbb{C}} \mathcal{M} = 2r$, $\dim_{\mathbb{C}} \mathcal{N} = 2r - 1$, and if C and D are \mathcal{M} -equivalent horizontal cylinders persisting on (Y, η) then they have identical heights and are not adjacent on (X, ω) .

Proof. By Corollary 5.4 the twist space of \mathcal{M} at (X, ω) is spanned by standard shears of \mathcal{M} -equivalence classes of cylinders. By Mirzakhani-Wright for cylinder collapses, there are disjoint collections of cylinder $\{\mathcal{C}_1, \dots, \mathcal{C}_n\}$ on (Y, η) such that the twist space of \mathcal{N} at (Y, η) is spanned by standard shears of $\mathcal{C}_1, \dots, \mathcal{C}_n$. Moreover, the degeneration theorem implies that two cylinders belong to the same \mathcal{C}_i if and only if they were \mathcal{M} -equivalent on (X, ω) . Therefore, the real dimension of the twist space of \mathcal{N} at (Y, η) is exactly the number of pairwise \mathcal{M} -inequivalent cylinders that persist on (Y, η) .

Suppose now that \mathcal{N} is higher rank. By Corollary 5.4 the twist space of (Y, η) at \mathcal{N} is spanned by standard shears of \mathcal{N} -equivalence classes. It follows immediately that $\mathcal{C}_1, \dots, \mathcal{C}_n$ are \mathcal{N} -equivalence classes of cylinders and so two cylinders are \mathcal{N} -equivalent if and only if their preimage on (X, ω) were \mathcal{M} -equivalent.

Suppose now that \mathcal{N} is rank one and that the twist space is two-dimensional, i.e. spanned by $u_{\mathcal{C}_1}$ and $u_{\mathcal{C}_2}$ where these vectors are standard shears and \mathcal{C}_1 and \mathcal{C}_2 are disjoint. By Theorem 5.5 it is immediate that \mathcal{N} is a branched covering construction of $\mathcal{H}(0, 0)$ and that any two \mathcal{M} -equivalent cylinders persisting on (Y, η) have identical heights and are not adjacent.

Suppose now that $\text{Twist}_{(X, \omega)} \mathcal{M} = \text{Pres}_{(X, \omega)} \mathcal{M}$. It follows from Theorem 5.1 and the results of Section 4 that (X, ω) has $\text{rk}(\mathcal{M}) + \text{rel}(\mathcal{M})$ equivalence classes of horizontal cylinders. If one cylinder from each equivalence class persists on (Y, η) then $\dim_{\mathbb{R}} \text{Twist}_{(Y, \eta)} \mathcal{N} = \text{rk}(\mathcal{M}) + \text{rel}(\mathcal{M})$. Since \mathcal{N} is contained in a hyperelliptic component $\text{rel}(\mathcal{N}) \leq 1$ and so $\dim_{\mathbb{C}} \mathcal{M} + (\text{rel}(\mathcal{M}) - 1) = 2 \cdot (\text{rk}(\mathcal{M}) + \text{rel}(\mathcal{M})) - 1 = 2 \cdot (\dim_{\mathbb{R}} \text{Twist}_{(Y, \eta)} \mathcal{N}) - 1 \leq \dim_{\mathbb{C}} \mathcal{N}$. By the degeneration theorem $\dim_{\mathbb{C}} \mathcal{N} < \dim_{\mathbb{C}} \mathcal{M}$ so $\text{rel}(\mathcal{M}) = 0$ and so $\dim_{\mathbb{C}} \mathcal{M} = 2r$ and $\dim_{\mathbb{C}} \mathcal{N} = \dim_{\mathbb{C}} \mathcal{M} - 1 = 2r - 1$. Now let C and D be \mathcal{M} -equivalent cylinders that persist on (Y, η) . If $\dim_{\mathbb{C}} \mathcal{N} = 3$, then the preceding proof shows that C and D are not adjacent and have identical heights. Otherwise, \mathcal{N} is a higher rank odd dimensional orbit closure in a hyperelliptic component of a stratum and $\text{Twist}_{(Y, \eta)} \mathcal{N} = \text{Pres}_{(Y, \eta)} \mathcal{N}$ since the dimension of the twist space is $\text{rk}(\mathcal{N}) + \text{rel}(\mathcal{N})$. Therefore, C and D are \mathcal{N} -equivalent on (Y, η) and hence have identical heights and are nonadjacent by Theorem 5.2. \square

8. BRANCHED COVERING CONSTRUCTIONS IN HYPERELLIPTIC COMPONENTS OF STRATA AND THE MIRZAKHANI CONJECTURE FOR ODD DIMENSIONAL ORBIT CLOSURES

Throughout this section \mathcal{M} will be a rank $r > 1$ affine invariant submanifold in a hyperelliptic component of a stratum. Suppose moreover that if $(X, \omega) \in \mathcal{M}$ is horizontally periodic and $\text{Twist}_{(X, \omega)} \mathcal{M} = \text{Pres}_{(X, \omega)} \mathcal{M}$ then

- (1) Any two \mathcal{M} -equivalent cylinders in (X, ω) have identical heights.

- (2) For any two adjacent cylinders in (X, ω) it is possible to (horizontally) shear their corresponding \mathcal{M} -equivalence classes to put the two cylinders in standard position.

By Lemma 4.8 it follows that any two \mathcal{M} -equivalent cylinders on any translation surface in \mathcal{M} have identical heights. Notice that condition 2 is trivial when the two cylinders lie in distinct \mathcal{M} -equivalence classes by the standard position lemma. Moreover, any odd dimensional higher rank affine invariant submanifold satisfies these hypotheses by the results of Section 5. Throughout this section we will suppose that (X, ω) is a horizontally periodic translation surface in \mathcal{M} such that $\text{Twist}_{(X, \omega)} \mathcal{M} = \text{Pres}_{(X, \omega)} \mathcal{M}$. The goal of the section is to show that \mathcal{M} is a branched covering construction by proving that (X, ω) is a simple translation covering.

Lemma 8.1. *(X, ω) contains a self-adjacent \mathcal{M} -equivalence class of horizontal cylinders if and only if $\dim_{\mathbb{C}} \mathcal{M} = 2r$.*

Proof. If $\dim_{\mathbb{C}} \mathcal{M} = 2r + 1$ then the results of Section 5 imply that since (X, ω) has $r + 1$ equivalence classes of horizontal cylinders no two \mathcal{M} -equivalent cylinders are adjacent.

Suppose to a contradiction that (X, ω) contains no self-adjacent \mathcal{M} -equivalence class of horizontal cylinders and $\dim_{\mathbb{C}} \mathcal{M} = 2r$. Since half-edges of the Lindsey tree produce self-adjacent cylinders it follows that the Lindsey tree has $g + 1$ vertices, i.e. (X, ω) has $g + 1$ horizontal cylinders. Since the Lindsey tree is a tree (not just a half-tree) there is a relation $\sum_c (-1)^{d(c, c_0)} \gamma_c = 0$ in homology where the sum is over cylinders, c_0 is an arbitrary fixed cylinder, and $d(c, c_0)$ is distance in the Lindsey tree. Since the tangent space of \mathcal{M} at (X, ω) has no relative deformations, we will reach a contradiction if we can show that $\sum_c (-1)^{d(c, c_0)} \gamma_c^*$ is in the tangent space of \mathcal{M} at (X, ω) . It suffices to show that \mathcal{M} -equivalent cylinders are an even distance apart in the Lindsey tree. This claim will follow from the following graph theory lemma:

Sublemma. *Let Γ be a tree whose vertices are colored with $\{1, \dots, m\}$. Suppose that*

- (1) *(No self-adjacency) No two vertices of the same color are adjacent.*
- (2) *(Cylinder Proportion Theorem) If v and w are vertices of the same color and v borders a vertex of color c then w does as well.*

Then vertices of the same color are an even distance apart in Γ .

Proof. Induct on the number of colors m . For $m = 2$ the result is clear. Now suppose that $m > 2$. Let a be a leaf of Γ connected to vertex b and suppose that a (resp. b) has color 1 (resp. 2). Let v and w be two vertices of the same color and let $[v, w]$ be the geodesic between them in Γ . Let Γ' be the colored tree that results from collapsing each maximal connected subtree of vertices of color 1 or 2 to a point of color 0. Let $[v]$ and $[w]$ be the images of v and w in Γ' . By the inductive hypothesis, $d([v], [w])$ is even. It suffices to show that $d(v, w)$ is even.

First suppose that v and w are not color 1. Let Z be the collection of points of color 0 on $[v, w]$ and for each $z \in Z$ let T_z be the maximal connected subtree of vertices of color 1 or 2 in Γ corresponding to z . Since all of the leaves of T_z are of color 2 it follows that the length of $[v, w] \cap T_z$, call it ℓ_z , is even since it is the length of a path between two vertices

of color 2 in T_z . Therefore, $d(v, w) = d([v], [w]) + \sum_{z \in Z} \ell_z$ is the sum of even numbers and hence even.

Now suppose v and w are of color 1. There are unique vertices of color 2, call them v' and w' , such that $d(v, w) = 2 + d(v', w')$. Since the distance between any two points of color 2 is even, it follows that the distance between v and w is even as well. \square

\square

Given a horizontally periodic translation surface with Lindsey tree Γ in a hyperelliptic component of a stratum, we will say that two edges or half-edges of Γ are \mathcal{M} -equivalent if they connect the same two \mathcal{M} -equivalence classes of horizontal cylinders. Half-edges will be understood to connect an \mathcal{M} -equivalence class to itself. If S is a collection of equivalent saddle connections on (X, ω) then let $S_t \cdot (X, \omega)$ be the translation surface where all saddle connections in S have been dilated by t .

Lemma 8.2. *If $\dim_{\mathbb{C}} \mathcal{M} = 2r$ then there is exactly one self-adjacent \mathcal{M} -equivalence class of horizontal cylinders on (X, ω) . If S is the collection of horizontal saddle connections that lie on the boundary of two \mathcal{M} -equivalent cylinders, then all saddle connections in S have identical lengths and $S_t \cdot (X, \omega) \in \mathcal{M}$.*

Proof. By lemma 8.1 there is at least one self-adjacent \mathcal{M} -equivalence class \mathcal{C}_0 of horizontal cylinders on (X, ω) . Let T be the collection of horizontal saddle connections that lie on the boundary of two cylinders in \mathcal{C}_0 .

If A and B are two adjacent cylinders in \mathcal{C}_0 then by assumption we may shear \mathcal{C}_0 so that A and B are in standard position with respect to one another. Assume without loss of generality that the saddle connection s on the boundary of A and B is the longest saddle connection in T . Let V be the vertical cylinder that forms between A and B and let \mathcal{V} be the \mathcal{M} -equivalence class of vertical cylinders that contains V . Horizontally collapse \mathcal{V} and let (Y, η) be the resulting translation surface on the boundary. By the vertical cylinder collapse lemma every component of (Y, η) is a translation surface that belongs to a hyperelliptic component of some stratum of abelian differentials.

Sublemma. *A cylinder from every \mathcal{M} -equivalence class persists on each component of (Y, η) .*

Proof. We begin by applying the cylinder deformation theorem to horizontally shrink \mathcal{V} so that any saddle connection completely contained in \mathcal{V} has length strictly smaller than the length of any other saddle connection. Now we will proceed by induction on the distance of an \mathcal{M} -equivalence class from \mathcal{C}_0 . The base case, that a cylinder from \mathcal{C}_0 persists on each component, is immediate. Suppose that there is an \mathcal{M} -equivalence class \mathcal{C}_1 at distance one from \mathcal{C}_0 . In particular this means that there is a cylinder C in \mathcal{C}_0 and a cylinder D in \mathcal{C}_1 so that C and D are adjacent. By the standard position lemma it is possible to put C and D into transverse standard position while fixing \mathcal{C}_0 . Let T be the resulting cylinder and let \mathcal{T} be the \mathcal{M} -equivalence class that contains T . By the cylinder proportion theorem for every cylinder in \mathcal{C}_0 there is a cylinder in \mathcal{T} that intersects it. Since the heights of all cylinders in \mathcal{T} are identical and strictly bigger than the length of any horizontal saddle connection contained in \mathcal{V} it follows that \mathcal{T} does not pass through any saddle connection

that is collapsed. It follows that a cylinder from \mathcal{C}_1 must persist on each component. Iterating this argument completes the proof. \square

The twist space degeneration lemma implies that adjacent \mathcal{M} -equivalent cylinders land on distinct components of (Y, η) and so \mathcal{V} contains all horizontal saddle connections in S . The cylinder proportion theorem implies that \mathcal{V} is contained in \mathcal{C}_0 since V is. Consequently, only \mathcal{C}_0 is self-adjacent. Since s is the longest saddle connection in S and is the width of V , it follows that every cylinder in \mathcal{V} has width that is identical in length to s . By the cylinder deformation theorem, dilating \mathcal{V} by t and then undoing the shear that put A and B in standard position remains in \mathcal{M} and establishes that $S_t \cdot (X, \omega) \in \mathcal{M}$. \square

Lemma 8.3 (Saddle Connection Dilation Lemma). *Let S be an equivalence class of horizontal saddle connections on (X, ω) , then every element of S has the same length and $S_t \cdot (X, \omega) \in \mathcal{M}$ for all t .*

Proof. Let \mathcal{C}_0 and \mathcal{C}_1 be two distinct adjacent \mathcal{M} -equivalence classes of horizontal cylinders on (X, ω) . Let S be the equivalence class of horizontal saddle connections connecting them. Let $s \in S$ have length ℓ and be the longest element of S . By lemma 8.2 suppose without loss of generality that any saddle connection connecting two \mathcal{M} -equivalent cylinders has length strictly smaller than ℓ . Suppose that s lies on the boundary of $C_0 \in \mathcal{C}_0$ and $C_1 \in \mathcal{C}_1$. By the standard position lemma suppose without loss of generality that C_0 and C_1 are in standard position and let V be the resulting cylinder. Let \mathcal{V} be the \mathcal{M} -equivalence class of vertical cylinders that contains V . Since every cylinder in \mathcal{V} has width ℓ it follows that no cylinder in \mathcal{V} passes through a horizontal saddle connection that connects two \mathcal{M} -equivalent cylinders.

Sublemma. *\mathcal{V} contains every element of S .*

Proof. Suppose to a contradiction that C is a cylinder in \mathcal{C}_0 and D is a cylinder in \mathcal{C}_1 so that \mathcal{V} does not contain the saddle connection connecting them. Horizontally collapse \mathcal{V} and let (Y, η) be the resulting boundary translation surface. By the horizontal collapse lemma (Y, η) is a disjoint union of connected translation surfaces in hyperelliptic components of strata. Let (Y', η') be the component of (Y, η) containing C and D . It is immediate from the cylinder proportion theorem and the fact that \mathcal{V} does not pass through any saddle connections joining two \mathcal{M} -equivalent cylinders that a cylinder from every \mathcal{M} -equivalence class persists on (Y', η') . By the twist space degeneration lemma, (Y', η') cannot contain the image of two adjacent \mathcal{M} -equivalent cylinders and $\dim_{\mathbb{C}} \mathcal{M} = 2r$. By Lemma 8.1 and the cylinder proportion theorem, (Y', η') must contain the image of two adjacent \mathcal{M} -equivalent cylinders, which is a contradiction. \square

By hypothesis, every element of \mathcal{V} is the same horizontal length across. Since s was chosen to be the longest saddle connection in S and since it is the width of V , it follows that all saddle connections in S have identical lengths. Since every element of S is contained in \mathcal{V} and since \mathcal{V} only passes through saddle connections in S , the cylinder deformation theorem implies that dilating \mathcal{V} horizontally by t for arbitrary t remains in \mathcal{M} . Now undoing the shears that put C_0 and C_1 in standard position remains in \mathcal{M} by the cylinder deformation theorem and is the translation surface $S_t \cdot (X, \omega)$ as desired. \square

Lemma 8.4. *Suppose that there are n balls arranged in a circle and that each ball has a color $\{1, \dots, m\}$. Suppose that every color appears at least once. If i and j are two balls, let (i, j) denote the collection of balls running clockwise between ball i and ball j not including balls i and j . Let $C_{(i,j)}$ be the multiset of colors contained in (i, j) . Suppose moreover that whenever $i \neq j$ and $k \neq \ell$ are all balls of the same color c and neither (i, j) nor (k, ℓ) contains balls of color c , then $C_{(i,j)} = C_{(k,\ell)}$. Then the cyclic order of the colors of the balls is periodic with period m .*

Proof. Let i and j be any two balls of the same color (without loss of generality 1) such that (i, j) has the fewest possible elements. If (i, j) is empty then it follows that all balls are the same color and $m = 1$. Suppose now that (i, j) is nonempty. By hypothesis it must contain every color in $\{2, \dots, m\}$. Since (i, j) has the fewest number of elements conditional on i and j having the same color, it follows that every ball must appear exactly once in (i, j) . If $n = m$, then we are done. Otherwise there is a ball k such that (j, k) is disjoint from (i, j) and contains no balls of color 1. It follows that $[i, k]$ has $2m$ elements and we may suppose without loss of generality that their colors are $(1, 2, \dots, m, 1, c_2, \dots, c_m)$ where (c_2, \dots, c_m) is a permutation of $(2, \dots, m)$. By assumption (i, j) has the fewest possible elements given that i and j have the same color. Therefore, $c_m = m$. This in turn implies that $c_{m-1} = m-1$ and so we have that the order of the colors on $[i, k]$ is $(1, \dots, m, 1, \dots, m)$. Iterating this argument gives that the order of the color of the balls is periodic with period m . \square

For each vertex v in the Lindsey tree of (X, ω) there is a cyclic ordering on edges attached to v given by the clockwise order of horizontal saddle connections on the top boundary of the cylinder corresponding to v . These edges are colored by associating to them the name of their equivalence class. When discussing the cyclic ordering around a vertex we will use “edge” to mean both edges and half-edges.

Lemma 8.5 (Periodic Ordering Lemma 1). *Let v be a vertex of the Lindsey tree. Label the equivalence classes of edges around v by $\{1, \dots, m\}$. The cyclic ordering of edge equivalence classes around v is (perhaps after relabelling) $(1, 2, \dots, m, 1, 2, \dots, m, \dots)$.*

Proof. By the Lemma 8.4 it suffices to show the following: suppose that i, j, k are saddle connections on the top boundary of v in equivalence class c . Suppose furthermore that (in the notation of Lemma 8.4) (i, j) and (j, k) do not contain saddle connections in equivalence class c . It suffices to show that $C_{(i,j)}$ and $C_{(j,k)}$ coincide as multi-sets.

Let a be any transcendental number and let $\{c_1, \dots, c_m\}$ be an enumeration of the edge equivalence classes around v . By the saddle connection dilation theorem we may suppose without loss of generality that each edge in equivalence class k has length a^k . By assumption we may suppose without loss of generality that j lies over $J(j)$ and is contained in a vertical cylinder V such that the only horizontal saddle connections intersected by V are j and $J(j)$. Let \mathcal{V} be the \mathcal{M} -equivalence class of vertical cylinders containing V . By assumption every cylinder in \mathcal{V} has the same width across. Therefore, \mathcal{V} only passes through horizontal saddle connections that are \mathcal{M} -equivalent to j . By the saddle connection dilation theorem, \mathcal{V} must contain every saddle connection equivalent to j in its union. In particular, i must lie above $J(k)$ and k must lie about $J(i)$. In other words the sum of the lengths of the saddle connections in (i, j) is identical to the sum of the lengths of the saddle connections in

(j, k) . Since we have assumed that \mathcal{M} -equivalence class k of horizontal saddle connections all have length a^k where a is transcendental, it follows that $C_{(i,j)} = C_{(k,\ell)}$. \square

Lemma 8.6 (Periodic Ordering Lemma 2). *Let v and w be \mathcal{M} -equivalent cylinders. Then the boundaries of v and w contain the same \mathcal{M} -equivalence classes of horizontal saddle connections in the same cyclic order.*

Proof. Suppose that the \mathcal{M} -equivalence classes of horizontal saddle connections adjacent to v are labelled by $\{0, \dots, m-1\}$ so that the cyclic ordering of horizontal saddle connections along the top boundary of v is $(0, 1, \dots, m-1, 0, 1, \dots, m-1, 0, \dots)$. Suppose first to a contradiction that \mathcal{M} -equivalence class 0 is adjacent to v but not to w . Applying the saddle connection dilation lemma we see that we may alter the length of the core curve of v while fixing the length of the core curve of w . This is impossible since Wright [22] implies that the ratio of lengths of core curves of \mathcal{M} -equivalent cylinders is always an algebraic number. Therefore, the boundaries of v and w contain the same \mathcal{M} -equivalence classes of horizontal saddle connections.

Apply the saddle connection dilation lemma to assume that the length of every horizontal saddle connection in equivalence classes $\{0, \dots, m-1\}$ has unit length. Apply the cylinder deformation theorem to assume without loss of generality that cylinders v and w have unit heights. If v borders an \mathcal{M} -equivalent cylinder then we will suppose that the corresponding edge equivalence class is labelled 0.

By the standard position lemma or perhaps by assumption we may put v into standard position with a cylinder v_0 in equivalence class 0. Suppose that the order of cylinders around the top boundary of v is $(v_0, \dots, v_{m-1}, v_m, \dots)$. By the standard position lemma fix v and put v into transverse standard position with v_1, \dots, v_{m-1} . Let \mathcal{T}_k be the equivalence classes of transverse cylinders that arise from putting v and v_k into transverse standard position. The slope of the core curve of cylinders in \mathcal{T}_k is $\frac{1}{2k}$. By the cylinder proportion theorem w contains a cylinder in \mathcal{T}_k for each k . It follows immediately that \mathcal{M} -equivalence classes of w appear in the same cyclic ordering as those in v . \square

Suppose that C_1, \dots, C_n are horizontal cylinders on a flat surface. Suppose that points on the boundary corresponding to zeros are marked. Say that C_1, \dots, C_n are mutually isogenous if there is a flat horizontal cylinder C with marked points on its boundary so that for each C_i there is a local isometry $f_i : C_i \rightarrow C$ so that the set of points sent to marked points of C are exactly the marked points of C_i .

Lemma 8.7 (Isogeny Lemma). *All \mathcal{M} -equivalent horizontal cylinders on (X, ω) are isogeneous.*

Proof. Let \mathcal{C} be an \mathcal{M} -equivalence class of horizontal cylinders. Let $C \in \mathcal{C}$ be a cylinder and let D be an adjacent \mathcal{M} -inequivalent cylinders. Put C and D into standard position by using the cylinder deformation theorem (suppose that \mathcal{C} is sheared by $\begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}$) - and let V be the resulting vertical cylinder. Let \mathcal{V} be the \mathcal{M} -equivalence class of cylinders containing V . Label the equivalence class containing the edge connecting C to D by 1. For any cylinder $v \in \mathcal{C}$ let $\{1, \dots, m\}$ be the edge equivalence classes attached to v and suppose that they are ordered by $(1, \dots, m, \dots)$ around v . By assumption every cylinder in \mathcal{C} has

Therefore all cylinders in \mathcal{C} are isogenous to $\begin{pmatrix} 1 & -t \\ 0 & 1 \end{pmatrix}$ applied to

$$h \begin{array}{|c|c|c|c|} \hline & \ell_1 & \ell_2 & \cdots \ell_m \\ \hline & \vdots & \vdots & \vdots \\ \hline & \ell_1 & \ell_m & \cdots \ell_2 \\ \hline \end{array}$$

Theorem 8.8 (Branched Cover Criteria). *Suppose that \mathcal{M} is a rank $r > 1$ affine invariant submanifold in a hyperelliptic component such that for any horizontally periodic (X, ω) with $\text{Twist}_{(X, \omega)} \mathcal{M} = \text{Pres}_{(X, \omega)} \mathcal{M}$*

- (1) *Any two \mathcal{M} -equivalent horizontal cylinders in (X, ω) have identical heights.*
- (2) *For any two adjacent horizontal cylinders in (X, ω) it is possible to (horizontally) shear their corresponding \mathcal{M} -equivalence classes to put the two cylinders in standard position.*

Then \mathcal{M} is a branched covering construction of $\mathcal{H}^{\text{hyp}}(2r-2)$ when $\text{rel}(\mathcal{M}) = 0$ and of $\mathcal{H}^{\text{hyp}}(r-1, r-1)$ when $\text{rel}(\mathcal{M}) = 1$. The cover is branched over zeros of the abelian differential and commutes with the hyperelliptic involution.

Proof. By Corollary 3.4 it suffices to show that if (X, ω) is a horizontally periodic translation surface with $\text{Twist}_{(X, \omega)} \mathcal{M} = \text{Pres}_{(X, \omega)} \mathcal{M}$ then (X, ω) is a simple translation covering of a generic surface (Y, η) in $\mathcal{H}^{hyp}(2r-2)$ or $\mathcal{H}^{hyp}(r-1, r-1)$. Enumerate the \mathcal{M} -equivalence classes of horizontal cylinders $\{1, \dots, m\}$ and the equivalence classes of edges $\{1, \dots, n\}$. Choose two transcendental numbers a and b so that $\mathbb{Q}(a, b)$ is isomorphic as a field to $\mathbb{Q}(x, y)$ where x and y are indeterminates. Using the cylinder deformation theorem and the saddle connection dilation lemma, assume without loss of generality that the height of the cylinders in equivalence class k is a^k and the lengths of saddle connections in equivalence class k is b^k .

We will first build the Lindsey tree Γ_Y of (Y, η) . For each \mathcal{M} -equivalence class in (X, ω) add a corresponding node in (Y, η) . If two distinct equivalence classes are adjacent in (X, ω) , then add an edge connecting the corresponding nodes in Γ_Y . If an equivalence class is self-adjacent in (X, ω) then add a half-edge to the corresponding node in Γ_Y . The cyclic order around each node is specified by the periodic ordering lemmas. This completely specifies Γ_Y .

We will now build (Y, η) and the simple translation covering $f : (X, \omega) \rightarrow (Y, \eta)$. To each \mathcal{M} -equivalence class \mathcal{C}_i of horizontal cylinders let C_i be the smallest cylinder isogenous to the \mathcal{C}_i (this will be the one constructed in Lemma 8.7). For each $c \in \mathcal{C}_i$ let $f_c : c \rightarrow C_i$ be the local isometry constructed in Lemma 8.7. Gluing the cylinders C_i together according to Γ_Y now constructs a translation surface (Y, η) . Let f be the map that sends a point $x \in c$ to $f_c(x)$. This map is conformal on the complement of $Z(\eta)$ - the zeros and marked points

of η . By the Riemann extension theorem f extends to a holomorphic map $f : X \rightarrow Y$ such that $f^*\eta = \omega$. In particular, $f : X - Z(\omega) \rightarrow Y - Z(\eta)$ is a degree d covering map.

Define a degree d branched covering of graphs to be a simplicial map between graphs $\pi : \Gamma \rightarrow \Gamma'$ such that

- (1) $|E_\Gamma| = d \cdot |E_{\Gamma'}|$ where E_Γ (resp. $E_{\Gamma'}$) is the edge set of Γ (resp. Γ').
- (2) For each $v \in \Gamma$ the ramification index $e_v := \frac{\deg(v)}{\deg \pi(v)}$ is an integer.
- (3) For each $w \in \Gamma'$, $\sum_{\pi(v)=w} (e_v - 1) = d$.

Heuristically the branch points will be a subset of the vertices. Riemann-Hurwitz for branched covers of graphs then says that

$$\chi_\Gamma = |V_\Gamma| - |E_\Gamma| = d \cdot |V_{\Gamma'}| - \sum_{v \in B} (e_v - 1) - d|E_{\Gamma'}| = d \cdot \chi_{\Gamma'} - \sum_{v \in B} (e_v - 1)$$

where B is the collection of branch points in Γ and χ denotes Euler characteristic.

We will now show that Γ_Y is a half-tree. Let Γ'_X and Γ'_Y be the graphs with half-edges deleted and, in (X, ω) , edges connecting two \mathcal{M} -equivalent cylinders deleted. Since $f : X - Z(\omega) \rightarrow Y - Z(\eta)$ is a degree d covering map it follows immediately that $\Gamma'_X \rightarrow \Gamma'_Y$ is a branched covering of graphs. Since $\chi_{\Gamma'_X}$ is positive and Γ'_Y is connected it follows that χ_{Γ_Y} is positive and hence 1, i.e. Γ_Y is a half-tree.

By Corollary 2.3, (Y, η) is a translation surface in a hyperelliptic component of a stratum. Moreover, the translation covering $f : (X, \omega) \rightarrow (Y, \eta)$ is a simple translation covering. Since Γ_Y is a half-tree on $\text{rk}(\mathcal{M}) + \text{rel}(\mathcal{M})$ vertices with an extra half-edge if and only if $\text{rel}(\mathcal{M}) = 0$, it follows that (Y, η) belongs to $\mathcal{H}^{\text{hyp}}(2r - 2)$ when $\text{rel}(\mathcal{M}) = 0$ and to $\mathcal{H}^{\text{hyp}}(r - 1, r - 1)$ when $\text{rel}(\mathcal{M}) = 1$. Moreover the map f is clearly branched over zeros and commutes with the hyperelliptic involution by construction.

We will now argue that \mathcal{M} is a branched covering construction. Since the heights and widths of inequivalent cylinders on (X, ω) were constructed from distinct transcendental numbers we have that the rational relations on moduli of cylinders are generated by $\frac{m_c}{m_{c'}} = q_{c, c'}$ for \mathcal{M} -equivalent horizontal cylinders c and c' and where m_c and $m_{c'}$ are moduli of cylinders c (resp. c') and $q_{c, c'}$ is a rational number. Therefore, the orbit closure of $\begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}$ contains the standard shears u_{c_i} for the $\text{rk}(\mathcal{M}) + \text{rel}(\mathcal{M})$ equivalence classes. It follows immediately that the orbit closure of (X, ω) is \mathcal{M} . Since (X, ω) is \mathcal{M} -generic and is a simple translation covering it follows by Corollary 3.4 that \mathcal{M} is a branched covering construction of \mathcal{H} where \mathcal{H} is the component of a stratum of abelian differentials containing (Y, η) . \square

As we have already observed the conditions in this theorem are satisfied by any odd dimensional affine invariant submanifold in a hyperelliptic component of a stratum.

Corollary 8.9. *Any odd-dimensional rank $r > 1$ affine invariant submanifold \mathcal{M} in a hyperelliptic component of a stratum is a branched covering construction of $\mathcal{H}^{\text{hyp}}(r - 1, r - 1)$.*

9. THE MIRZAKHANI CONJECTURES: EVEN DIMENSIONAL ORBIT CLOSURES

Throughout this section \mathcal{M} will be an affine invariant submanifold in $\mathcal{H}^{hyp}(2g-2)$ or $\mathcal{H}^{hyp}(g-1, g-1)$ of even complex dimension $2r$ where $r > 1$. The goal of this section is to show that \mathcal{M} is a branched covering construction over $\mathcal{H}^{hyp}(2r-2)$.

Theorem 9.1. *If $r = 2$ then \mathcal{M} is a branched covering construction of $\mathcal{H}(2)$.*

Proof. By Wright [22] there is a translation surface in \mathcal{M} with two distinct \mathcal{M} -equivalence classes of horizontal cylinders. By the perturbation theorem we may replace this translation surface with (X, ω) - a new translation surface that contains an \mathcal{M} -equivalence class of horizontal cylinders \mathcal{C}_0 and a disjoint vertical \mathcal{M} -equivalence class of vertical cylinders \mathcal{C}_1 . Applying the cylinder deformation we may assume without loss of generality that \mathcal{C}_0 contains a vertical saddle connection and that \mathcal{C}_1 contains a horizontal saddle connection. By Mirzakhani-Wright [12] collapsing \mathcal{C}_1 passes to a disjoint union of translation surfaces (Y, η) on the boundary of \mathcal{M} , each component of which has orbit closure whose dimension is less than or equal to three. In particular, each component is completely periodic. Let (Y', η') denote the boundary translation surface formed by collapsing \mathcal{C}_0 .

Lemma. \mathcal{C}_0 appears on every component of (Y, η) .

Proof. Let C be a connected component of $(X, \omega) - \mathcal{C}_0 - \mathcal{C}_1$. Suppose to a contradiction that C does not border \mathcal{C}_0 . It follows that C borders a cylinder D in \mathcal{C}_1 and so C and D land on the same component (Y'_0, η'_0) of (Y', η') . Since (Y'_0, η'_0) is completely periodic and contains vertical cylinder D it follows that C is a union of vertical cylinders. Since C does not border \mathcal{C}_0 these vertical cylinders persist on (X, ω) and cover C .

By assumption the vertical cylinders in C are not \mathcal{M} -equivalent to the cylinders in \mathcal{C}_1 . Let the equivalence class of the cylinders in C be called \mathcal{C}_2 . Since \mathcal{M} has dimension 4 it follows that \mathcal{C}_1 and \mathcal{C}_2 cover (X, ω) . However, the cylinder proportion theorem implies that both \mathcal{C}_1 and \mathcal{C}_2 are disjoint from \mathcal{C}_0 , which is a contradiction. \square

Since each component of (Y, η) is completely periodic and contains a horizontal cylinder from \mathcal{C}_0 it follows that (Y, η) is horizontally periodic. Let \mathcal{D}_0 be the second equivalence class of horizontal cylinders. Similarly, let \mathcal{D}_1 be the second equivalence class of vertical cylinders on each component of (Y', η') .

Lemma. No two cylinders in \mathcal{C}_0 are adjacent.

Proof. Suppose to a contradiction that two cylinders C and D in \mathcal{C}_0 are adjacent. Let (Y_0, η_0) be the component of (Y, η) containing C and D . Let \mathcal{N} be the orbit closure of (Y_0, η_0) . Mirzakhani-Wright implies that the twist space of (Y_0, η_0) contains the standard shear of horizontal cylinders in \mathcal{C}_0 . It follows that the twist space also contains the standard shear of horizontal cylinders in \mathcal{D}_0 . Let $a_t^{\mathcal{D}_0}$ and $a_t^{\mathcal{C}_0}$ be vertical dilation of \mathcal{D} (resp. \mathcal{C}_0) by t . Since the standard shear of \mathcal{C}_0 and \mathcal{D}_0 are in the twist space, $a_t^{\mathcal{D}_0}$ and $a_t^{\mathcal{C}_0}$ are in $T_{(Y_0, \eta_0)}\mathcal{N}$.

Since $\dim_{\mathbb{C}} \mathcal{N} \leq 3$ it follows that $\text{rk}(\mathcal{N}) = 1$. Consequently some linear combination α of $a_1^{\mathcal{D}}$ and $a_1^{\mathcal{C}_0}$ is a relative deformation. Suppose first to a contradiction that the coefficient of $a_1^{\mathcal{C}_0}$ in α is nonzero. After scaling $\alpha = a_1^{\mathcal{C}_0} + ca_1^{\mathcal{D}_0}$ for some $c \in \mathbb{C}$. However, this cannot be a relative deformation since $C \cup D$ contains an absolute period and α increases the length of this path, a contradiction. Therefore, after scaling, $\alpha = a_1^{\mathcal{D}_0}$. But this deformation increases the length of every path and hence cannot be a relative deformation. We have a contradiction. \square

The horizontal and vertical collapse lemmas imply that (Y, η) and (Y', η') are disjoint unions of translation surfaces in hyperelliptic components of strata. Moreover, (Y, η) (resp. (Y', η')) is horizontally (resp. vertically) periodic. It is immediate that (X, ω) is vertically and horizontally periodic. By abuse of notation, let \mathcal{D}_0 (resp. \mathcal{D}_1) be the \mathcal{M} -equivalence class of horizontal (resp. vertical) cylinders on (X, ω) that projects to \mathcal{D}_0 (resp. \mathcal{D}_1) on (Y, η) (resp. (Y', η')). By the twist space degeneration lemma cylinders in \mathcal{D}_0 are not adjacent on (Y, η) and cylinders in \mathcal{D}_1 are not adjacent on (Y', η') . Therefore, \mathcal{C}_0 (resp. \mathcal{C}_1) contains all saddle connections on the boundary of two cylinders in \mathcal{D}_1 (resp. \mathcal{D}_0) and these are exactly the saddle connections that these cylinders intersect. Finally, the twist space degeneration lemma implies that if two \mathcal{M} -equivalent horizontal or vertical cylinders land on the same component of a degeneration then they have the same height.

Lemma. *All cylinders in \mathcal{C}_0 have identical heights.*

Proof. Given two horizontal cylinders C and D let (C, D) be the geodesic between them in the Lindsey tree. We have already seen that if C and D are two cylinders in \mathcal{C}_0 then there are cylinders $C_i \in \mathcal{C}_0$ for $0 \leq i \leq n$ such that $C_0 = C$, $C_n = D$, and (C_i, C_{i+1}) intersects at most one cylinder in \mathcal{C}_1 . If (C_i, C_{i+1}) intersects no cylinder in \mathcal{C}_1 then C_i and C_{i+1} land on same component of (Y, η) and hence have the same height. Therefore to show that all cylinders in \mathcal{C}_0 have identical heights it suffices to show that if C and D are cylinders in \mathcal{C}_0 and (C, D) intersects \mathcal{C}_1 exactly once that C and D have identical heights.

Let V be the unique cylinder in \mathcal{C}_1 that (C, D) enters. Let s (resp. s') be the vertical saddle connection on the left (resp. right) boundary of V through which (C, D) passes. By the cylinder deformation theorem it is possible to shear \mathcal{C}_1 so that V contains a horizontal saddle connection and so s and s' are connected by a horizontal line in V . Collapse \mathcal{C}_1 . Now C and D land on the same component of the boundary translation surface and the previous lemmas imply that C and D have identical heights. \square

Lemma. *All cylinders in \mathcal{D}_0 have identical heights.*

Proof. Let D be the shortest cylinder in \mathcal{D}_0 bordering a cylinder in \mathcal{C}_0 . Let C be a cylinder in \mathcal{C}_0 that D borders. By the standard position lemma put C and D into standard position and let V be the resulting vertical cylinder. Let \mathcal{V} be the equivalence class of vertical cylinders containing V . Every cylinder in \mathcal{V} must spend the same percent of time in \mathcal{D}_0 that V does. Since every cylinder in \mathcal{C}_0 has the same height and no two are adjacent, this implies that every cylinder in \mathcal{D}_0 has the same height as D . \square

Lemma. *All cylinders in \mathcal{C}_1 and \mathcal{D}_1 have identical widths by symmetry of hypotheses.*

In summary, (X, ω) is horizontally and vertically periodic and \mathcal{M} -equivalent horizontal (resp. vertical) cylinders have identical heights (resp. widths). Equivalent edges have identical lengths. Since (Y', η') is a torus cover it also follows that equivalence classes of edges are periodically ordered with period two around each cylinder in \mathcal{D}_0 . The proof of Theorem 8.8 shows that (X, ω) is a simple translation covering of a translation surface in $\mathcal{H}(2)$. By corollary 3.4 \mathcal{M} is a branched covering construction over $\mathcal{H}(2)$ as desired. \square

Lemma 9.2 (Leaf Lemma). *Suppose that $(X, \omega) \in \mathcal{M}$ has r \mathcal{M} -equivalence classes of horizontal cylinders, then there is a horizontal cylinder that is adjacent to at most itself and an \mathcal{M} -inequivalent cylinder.*

Proof. Let Γ be the Lindsey half-tree of (X, ω) and suppose that each vertex is colored based on which \mathcal{M} -equivalence class the corresponding cylinder belongs to. Let Γ' be the quotient tree of Γ where each monochromatic subtree is collapsed to a point and half-edges are deleted. Let λ be a leaf of Γ' and let T be the corresponding monochromatic subtree in Γ . There is a cylinder $v \in T$ and an inequivalent cylinder w so that T is connected to the rest of the graph by the edge (v, w) . By the standard position lemma we may suppose without loss of generality that v and w are in standard position. However, this implies that the only cylinder in T that lies in a vertical cylinder connected to a cylinder equivalent to w is v . Therefore, $T = \{v\}$. Consequently, v is adjacent to at most itself and an \mathcal{M} -inequivalent cylinder w . \square

Lemma 9.3 (Identical Heights Lemma). *Suppose that $(X, \omega) \in \mathcal{M}$ has r \mathcal{M} -equivalence classes of horizontal cylinders and suppose that at most one of them is self-adjacent. Suppose that \mathcal{C} is the unique self-adjacent equivalence class (provided that one exists) and any equivalence class if no self-adjacent ones exist. If any two cylinders in \mathcal{C} have identical heights, then any two \mathcal{M} -equivalent horizontal cylinders on (X, ω) have identical heights.*

Proof. Induct on distance d from \mathcal{C} . The $d = 0$ base case holds by hypothesis. Suppose that \mathcal{D} is an \mathcal{M} -equivalence class of horizontal cylinders at distance $d > 0$ from \mathcal{C} . By assumption \mathcal{D} is not self-adjacent. Let \mathcal{D}' be an \mathcal{M} -equivalence class that is adjacent to \mathcal{D} that lies at distance $d - 1$ from \mathcal{C} . Let C be the tallest cylinder in \mathcal{D} and suppose that it is adjacent to $D \in \mathcal{D}'$. By the standard position lemma it is possible to put C and D into standard position. Let W be the resulting cylinder and let \mathcal{W} be the \mathcal{M} -equivalence class of vertical cylinders containing W . Since \mathcal{D} is not self-adjacent and every cylinder in \mathcal{D}' has the same height, the cylinder proportion theorem implies that every cylinder in \mathcal{D} has the same height as C as desired. \square

Theorem 9.4. \mathcal{M} is a branched covering construction over $\mathcal{H}^{hyp}(2r - 2)$.

Proof. Proceed by induction on r . The base case is Theorem 9.1. Let $(X, \omega) \in \mathcal{M}$ have r \mathcal{M} -equivalence classes of horizontal cylinders and let v be a cylinder produced by the leaf lemma. Let \mathcal{C}_0 be the \mathcal{M} -equivalence class containing v . Let w be the unique \mathcal{M} -inequivalent cylinder adjacent to v . Let \mathcal{C}_1 be the \mathcal{M} -equivalence class of w .

Case 1: v has a half-edge in the Lindsey tree

Suppose without loss of generality, perhaps after using the cylinder deformation theorem to shear \mathcal{C}_0 , that the half-edge corresponds to a vertical cylinder V contained in v . Let \mathcal{V} be the \mathcal{M} -equivalence class of vertical cylinders containing V . Let (Y, η) be the vertical cylinder collapse of \mathcal{V} . By the horizontal collapse lemma, (Y, η) is a disjoint union of translation surfaces in hyperelliptic components of strata.

Lemma. *Each component of (Y, η) either contains only cylinders that belonged to \mathcal{C}_0 or they contain a cylinder from every \mathcal{M} -equivalence class of horizontal cylinders.*

Proof. First we will slightly alter (X, ω) . Proceed recursively. An \mathcal{M} -equivalence class \mathcal{C} at distance d from \mathcal{C}_0 is adjacent to some \mathcal{M} -equivalence class \mathcal{C}' of distance $d - 1$. Put \mathcal{C} and \mathcal{C}' in transverse standard position while fixing \mathcal{C}' . These transverse cylinders persist as \mathcal{V} is shrunk and therefore the cylinder proportion theorem implies that any component that contains a cylinder from \mathcal{C}_1 contains every \mathcal{M} -equivalence class of horizontal cylinders. \square

The twist space degeneration lemma implies that \mathcal{C}_0 is the only self-adjacent \mathcal{M} -equivalence class of horizontal cylinders and that if C and D are two \mathcal{M} -equivalent cylinders not contained in \mathcal{C}_0 that land on the same component of (Y, η) then they have identical heights.

Since $r > 2$ there is an \mathcal{M} -equivalence class of horizontal cylinders \mathcal{C}_2 that is not adjacent to \mathcal{C}_0 . We may suppose without loss of generality, perhaps after applying the cylinder deformation theorem, that \mathcal{C}_2 contains a vertical saddle connection. Vertically collapse \mathcal{C}_2 . By the horizontal cylinder collapse lemma the resulting translation surface (Y', η') is a disjoint union of translation surfaces in a hyperelliptic component of a stratum. Any component of (Y', η') that contains a cylinder from \mathcal{C}_0 also contains a cylinder from \mathcal{C}_1 . Since \mathcal{C}_0 is self-adjacent it follows that the orbit closures of these components cannot have rank one. The twist space degeneration lemma then implies that cylinders in \mathcal{C}_0 that persist on the same component of (Y', η') remain equivalent there. The induction hypothesis implies that \mathcal{C}_0 may be sheared on that component to put any two adjacent cylinders in \mathcal{C}_0 in standard position. The shear may be performed on (X, ω) to put the two cylinders in standard position.

Lemma. *All cylinders in \mathcal{C}_1 have identical heights.*

Proof. Any path in the Lindsey tree between two vertices in \mathcal{C}_1 is a concatenation of paths between vertices in \mathcal{C}_1 so that the path either always or never passes through vertices contained in \mathcal{C}_0 . In the first case the endpoints of the subpath persist on the same component of (Y', η') and in the second the endpoints of the path persist on the same component of (Y, η) . As remarked above this implies that the two endpoints of each subpath have identical heights. Therefore any two cylinders in \mathcal{C}_1 have identical heights. \square

Let C be the shortest cylinder in \mathcal{C}_0 adjacent to a cylinder in \mathcal{C}_1 . Let D be the cylinder to which C is adjacent. Put C and D in standard position. Since \mathcal{C}_1 is not self-adjacent and all cylinders in it have identical heights, the cylinder proportion theorem implies that the cylinders in \mathcal{C}_0 have identical heights. The identical heights lemma now implies that any two \mathcal{M} -equivalent horizontal cylinders on (X, ω) have identical heights. Therefore, the branched covering criteria hold in this case.

Case 2: v does not have a half-edge

By the standard position lemma we may suppose without loss of generality that v and w are in standard position and \mathcal{W} is the resulting \mathcal{M} -equivalence class of vertical cylinders.

Lemma. *\mathcal{C}_0 is not self-adjacent.*

Proof. Suppose not to a contradiction and let C and C' be adjacent (perhaps identical) cylinders contained in \mathcal{C}_0 . Since $r > 2$ there is an \mathcal{M} -equivalence class of horizontal cylinders \mathcal{C}_2 that is not adjacent to \mathcal{C}_0 .

If \mathcal{C}_2 is not self-adjacent then we will use the cylinder deformation theorem to shear \mathcal{C}_2 so that it contains a vertical saddle connection, vertically collapse, and thereby pass to a disjoint union of translation surfaces in hyperelliptic components of strata (Z, ζ) by the horizontal cylinder collapse lemma. The component of (Z, ζ) that contains C and C' has orbit closure whose twist space is at least two-dimensional. By the induction hypothesis and the main theorem for odd dimensional orbit closures it follows that C and C' may be put in standard position by shearing \mathcal{C}_0 . This must remain true on (X, ω) .

If \mathcal{C}_2 is self-adjacent then let D and D' be adjacent (perhaps identical) cylinders contained in \mathcal{C}_2 . Horizontally collapse \mathcal{W} to pass to a disjoint union of translation surfaces in hyperelliptic components of strata (W, ξ) by lemma 7.1. The orbit closure of the component of (W, ξ) that contains D and D' must have orbit closure whose twist space is at least two-dimensional. By the induction hypothesis and the main theorem for odd dimensional orbit closures it follows that D and D' may be put in standard position by shearing \mathcal{C}_2 . This must remain true on (X, ω) . Let V be the resulting vertical cylinder. Let \mathcal{V} be the \mathcal{M} -equivalence class containing V . Now horizontally collapse \mathcal{V} . By the vertical cylinder collapse lemma this passes to a translation surface (Z, ζ) that is a disjoint union of translation surfaces in hyperelliptic components of strata. By the induction hypothesis and the main theorem for odd dimensional orbit closures it follows that C and C' may be put in standard position by shearing \mathcal{C}_0 . This must remain true on (X, ω) .

Let U be the vertical cylinder that forms between C and C' now that they are in standard position. The cylinder proportion theorem implies that there is an \mathcal{M} -equivalent vertical cylinder contained in v . However this is impossible since v is a leaf without half-edges. \square

Vertically collapse \mathcal{C}_0 . By the horizontal cylinder collapse lemma the resulting translation surface (Y, η) is a disjoint union of translation surfaces in hyperelliptic components of strata. Moreover, all \mathcal{M} -equivalence classes, except \mathcal{C}_0 , persist on each component of (Y, η) . By the induction hypothesis, Theorem 5.5, and Theorem 8.9 two \mathcal{M} -equivalent cylinders that land on the same component of (Y, η) have identical heights. Since \mathcal{C}_0 is not self-adjacent, the induction hypothesis, Theorem 5.5, and Theorem 8.9 imply that for any two adjacent cylinders it is possible to shear their equivalence class to put them in standard position.

Lemma. *All cylinders in \mathcal{C}_1 have identical heights.*

Proof. Recall that (Z, ζ) was the boundary translation surface constructed above that results either from collapsing \mathcal{C}_2 if \mathcal{C}_2 is not self-adjacent or collapsing an \mathcal{M} -equivalence class of vertical cylinders contained in \mathcal{C}_2 otherwise. Let C and C' be cylinders in \mathcal{C}_1 . If the geodesic between them in the Lindsey tree only passes through cylinders in \mathcal{C}_0 , then C and C' land on the same component of (Z, ζ) and hence have identical heights. If the geodesic passes through no cylinder in \mathcal{C}_0 then C and C' land on the same component of (Y, η) and hence have the same height. Since the geodesic between C and C' in the Lindsey tree is a concatenation of paths between cylinders in \mathcal{C}_1 that either avoid or exclusively pass through \mathcal{C}_0 , it follows that C and C' have identical heights. \square

Lemma. *The argument in the identical heights lemma shows that all cylinders in \mathcal{C}_0 have identical heights.*

If \mathcal{C}_1 is the only self-adjacent \mathcal{M} -equivalence class or if no \mathcal{M} -equivalence class of horizontal cylinders is self-adjacent then the second half of the branched cover criteria holds by the identical heights lemma. Therefore suppose that there is some \mathcal{M} -equivalence class \mathcal{C} of horizontal cylinders distinct from \mathcal{C}_0 and \mathcal{C}_1 that is self-adjacent.

Lemma. *\mathcal{C} is the unique self-adjacent \mathcal{M} -equivalence class on (X, ω) .*

Proof. Recall that (Y, η) is the horizontal cylinder collapse of \mathcal{C}_0 and that \mathcal{W} is the \mathcal{M} -equivalence class of vertical cylinders that contains W - the vertical cylinder that was

created by putting vertices v and w in standard position. As in the first case a cylinder from every \mathcal{M} -equivalence class of horizontal cylinders, except \mathcal{C}_0 , persists on each component of (Y, η) . The twist space degeneration lemma implies that the dimension of the twist space of the orbit closure of any component of (Y, η) is $r - 1$ and that two cylinders that persist on the same component of (Y, η) are equivalent on the orbit closure if and only if they were \mathcal{M} -equivalent on (X, ω) . Collapse \mathcal{W} on each component of (Y, η) and let (Y'', η'') be the resulting translation surface. By the vertical cylinder collapse lemma (Y'', η'') is a disjoint union of translation surfaces in hyperelliptic components of strata. Since a cylinder from each \mathcal{M} -equivalence class continues to persist on each component of (Y'', η'') the twist space degeneration lemma implies that the dimension of the twist space of the orbit closure of any component of (Y'', η'') is at least $r - 1$ and hence each component of (Y'', η'') has orbit closure of dimension either $2r - 2$ or $2r - 3$. If the dimension is $2r - 3$ then since the twist space is maximal dimension it follows that no two \mathcal{M} -equivalent cylinders remain adjacent on any component of (Y'', η'') by the twist space degeneration lemma; this contradicts the fact that cylinders in \mathcal{C} remain adjacent on all components of (Y'', η'') .

Since $r \geq 3$ it follows that the orbit closure of each component of (Y'', η'') has dimension at least 4 and hence is higher rank. By the twist space degeneration lemma \mathcal{M} -equivalence classes of horizontal cylinders persisting on the same component of (Y'', η'') remain equivalent. By the induction hypothesis there is a unique self-adjacent equivalence class on each component of (Y'', η'') and this equivalence class must be \mathcal{C} . \square

Lemma. *Any two cylinders in \mathcal{C} have identical heights.*

Proof. Put two adjacent cylinders in \mathcal{C} in standard position and let \mathcal{V} be the resulting equivalence class of vertical cylinders. Collapse \mathcal{V} . By the vertical cylinder collapse lemma the resulting translation surface (Y''', η''') is a disjoint union of translation surfaces in hyperelliptic components of strata. By the induction hypothesis, Theorem 5.5, and Theorem 8.9 two \mathcal{M} -equivalent cylinders that land on the same component of (Y''', η''') have identical heights.

Let C and C' be cylinders in \mathcal{C} such that the geodesic between them in the Lindsey tree does not pass through any other cylinders contained in \mathcal{C} . It suffices to show that C and C' have identical heights. If C and C' are adjacent then they land on the same component of (Y, η) and hence have the same height. If C and C' are not adjacent then they land on the same component of (Y''', η''') and hence have the same height. \square

The identical heights lemma now establishes that any two \mathcal{M} -equivalent cylinders on (X, ω) have identical heights. The branched covering criteria now implies that \mathcal{M} is a branched covering construction of $\mathcal{H}^{hyp}(2r - 2)$ as desired. \square

REFERENCES

- [1] David Auricino and Duc-Manh Nguyen. Rank two affine submanifolds in $\mathcal{H}(2, 2)$ and $\mathcal{H}(3, 1)$. preprint, arXiv:1501.03303 (2015).
- [2] David Auricino, Duc-Manh Nguyen, and Alex Wright. Classification of higher rank orbit closures in $\mathcal{H}^{odd}(4)$. preprint, arXiv:1308.5879 (2013).

- [3] Artur Avila, Alex Eskin, and Martin Möller. Symplectic and Isometric $SL(2, \mathbb{R})$ -invariant subbundles of the Hodge bundle. preprint, arXiv 1209.2854 (2012).
- [4] Matt Bainbridge, Philipp Habegger, and Martin Moeller. Teichmueller curves in genus three and just likely intersections in $G_m^n \times G_a^n$. preprint, arXiv:1410.6835 (2014).
- [5] Kariane Calta. Veech surfaces and complete periodicity in genus two. *J. Amer. Math. Soc.*, 17(4):871–908, 2004.
- [6] Alex Eskin and Maryam Mirzakhani. Invariant and stationary measures for the $SL(2, \mathbb{R})$ action on moduli space. preprint, arXiv 1302.3320 (2013).
- [7] Alex Eskin, Maryam Mirzakhani, and Amir Mohammadi. Isolation theorems for $SL(2, \mathbb{R})$ -invariant submanifolds in moduli space. preprint, arXiv:1305.3015 (2013).
- [8] Maxim Kontsevich and Anton Zorich. Connected components of the moduli spaces of Abelian differentials with prescribed singularities. *Invent. Math.*, 153(3):631–678, 2003.
- [9] Erwan Lanneau, Duc-Manh Nguyen, and Alex Wright. Finiteness of Teichmüller curves in non-arithmetic rank 1 orbit closures. preprint, arXiv:1504.03742 (2015).
- [10] Kathryn Lindsey. Counting invariant components of hyperelliptic translation surfaces. preprint, arXiv 1302.3282. (2013).
- [11] Bernard Maskit. Comparison of hyperbolic and extremal lengths. *Ann. Acad. Sci. Fenn. Ser. A I Math.*, 10:381–386, 1985.
- [12] Carlos Matheus and Alex Wright. Hodge-Teichmüller planes and finiteness results for Teichmüller curves. *Duke Math. J.*, 164(6):1041–1077, 2015.
- [13] Curtis T. McMullen. Billiards and Teichmüller curves on Hilbert modular surfaces. *J. Amer. Math. Soc.*, 16(4):857–885 (electronic), 2003.
- [14] Curtis T. McMullen. Teichmüller curves in genus two: torsion divisors and ratios of sines. *Invent. Math.*, 165(3):651–672, 2006.
- [15] Curtis T. McMullen. Dynamics of $SL_2(\mathbb{R})$ over moduli space in genus two. *Ann. of Math. (2)*, 165(2):397–456, 2007.
- [16] Curtis T. McMullen. Navigating moduli space with complex twists. *J. Eur. Math. Soc. (JEMS)*, 15(4):1223–1243, 2013.
- [17] Maryam Mirzakhani and Alex Wright. The boundary of an affine invariant submanifold. preprint, arXiv:1508.01446 (2015).
- [18] Martin Möller. Finiteness results for Teichmüller curves. *Ann. Inst. Fourier (Grenoble)*, 58(1):63–83, 2008.
- [19] Duc-Manh Nguyen and Alex Wright. Non-Veech surfaces in $\mathcal{H}^{\text{hyp}}(4)$ are generic. *Geom. Funct. Anal.*, 24(4):1316–1335, 2014.
- [20] John Smillie and Barak Weiss. Characterizations of lattice surfaces. *Invent. Math.*, 180(3):535–557, 2010.
- [21] Alex Wright. The field of definition of affine invariant submanifolds of the moduli space of abelian differentials. *Geom. Topol.*, 18(3):1323–1341, 2014.
- [22] Alex Wright. Cylinder deformations in orbit closures of translation surfaces. *Geom. Topol.*, 19(1):413–438, 2015.
- [23] Alex Wright. Translation surfaces and their orbit closures: An introduction for a broad audience. *EMS Surv. Math. Sci.*, 2(1):63–108, 2015.
- [24] Anton Zorich. Flat surfaces. In *Frontiers in number theory, physics, and geometry. I*, pages 437–583. Springer, Berlin, 2006.